

# THINKING AND WRITING

## *Cognitive Science and Intelligence Analysis*

By Robert S. Sinclair



Revised edition of a monograph originally published by CSI in January 1984.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>FEB 2010</b>		2. REPORT TYPE		3. DATES COVERED <b>00-00-2010 to 00-00-2010</b>	
4. TITLE AND SUBTITLE <b>Thinking and Writing: Cognitive Science and Intelligence Analysis</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Central Intelligence Agency, Center for the Study of Intelligence, Washington, DC, 20505</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>41</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

**All statements of fact, opinion, or analysis expressed in this study are those of the author. They do not necessarily reflect official positions or views of the Central Intelligence Agency or any other US Government entity, past or present. Nothing in the contents should be construed as asserting or implying US Government endorsement of the study's factual statements and interpretations.**

Cover image creator, Bruno Mallart.  
© Images.com/Corbis.

---

# **Thinking and Writing: *Cognitive Science and Intelligence Analysis***

**Robert S. Sinclair**



***Center for the Study of Intelligence***

**Washington, DC**

**February 2010**

**Originally published in January 1984**

---

---

This study was originally prepared by Robert Sinclair during a fellowship with the Center for the Study of Intelligence. It has been updated with a new introduction and slightly edited.

Robert Sinclair was a CIA intelligence analysts for 37 years. He is now a consultant on analytical issues.

---

The Center for the Study of Intelligence (CSI) was founded in 1974 in response to Director of Central Intelligence James Schlesinger's desire to create within CIA an organization that could "think through the functions of intelligence and bring the best intellects available to bear on intelligence problems." The Center, comprising professional historians and experienced practitioners, attempts to document lessons from past operations, explore the needs and expectations of intelligence consumers, and stimulate serious debate on current and future intelligence challenges.

To support these activities, CSI publishes *Studies in Intelligence*, as well as books and monographs addressing historical, operational, doctrinal, and theoretical aspects of the intelligence profession. It also administers the CIA Museum and maintains the Agency's Historical Intelligence Collection.

Readers wishing to have copies of this monograph or other CSI publications may call 703-613-1753 or secure 31753. Requests may also be sent via Lotus Notes to CSI-PubReq.

---

---

# Contents

<b><i>Introduction and Update</i></b>	v
A Quick Summary	vi
Suggestion: Develop Collaborative Analysis	vii
Suggestion: Increase Cognitive Diversity	ix
 <b><i>Thinking and Writing: Cognitive Science and Intelligence Analysis</i></b>	
Summary	1
Recommendations	3
Preface	4
Introduction	4
I. A Quick Survey	6
II. Heuristics and the Incredible Chunk	8
III. Speaking and Writing, News and Knowledge	14
IV. Writing Schemes and Cognitive Overload	19
V. Creativity and the Conceptual Front End	21
VI. What Next?	26
1. Recruitment.	26
2. Training.	26
3. Analytic Writing.	26
4. Organization.	26
5. Presentational Methods.	27
6. Further Exploration.	27
Afterword	28

## Introduction and Update

When this monograph was published a quarter-century ago, it sank virtually without a trace. It is clear to me now that the paper lacked what today would be called “curb appeal”; moreover, cognitive science was a new and unproven discipline. Then, few inside or outside the intelligence world were aware of it, and even fewer had thought about its relevance to intelligence analysis.<sup>a</sup>

The field has opened up to a stunning degree since then. Not only have we seen a flood of studies documenting the myriad cognitive activities our brains engage in, but electronic imaging allows us to observe what happens in the brain as it goes about its business. Authors like Malcolm Gladwell have mined the literature to show the insights these processes can produce, as well as the times they leave us stuck in unproductive ways of thinking. In *Blink: The Power of Thinking Without Thinking*, Gladwell reports that when experts were asked to assess the provenance of an allegedly ancient sculpture, they could agree that it was a fake but could not put into words how they had reached that

conclusion. This and other examples, he says, illustrate how the things we learn through experience often are not readily available to our conscious minds.<sup>b</sup>

Economists these days speak of “behavioral economics,” which uses research based partly on cognitive-science protocols to suggest the limits to rational-actor models. Behavioral economics has reinstalled John Maynard Keynes on his pedestal.

*In economics, the crucial Keynesian concept is uncertainty. Where it prevails, the simple rules of classical economics don't apply. That's because the classical economics that both predated Keynes and superseded him relies on rational actors making rational assessments. In order to make such assessments you have to have reliable knowledge, usually derived from past experience. Buyers of oranges or newspapers or legal services can be said to possess such knowledge. Buyers of specula-*

---

a. One exception was Richards J. Heuer, whose articles, published in *Studies in Intelligence* beginning in 1978, helped trigger my own interest in the field. The CIA Center for the Study of Intelligence published an updated version of Heuer's articles in a book, *Psychology of Intelligence Analysis*, in 1999. Since reprinted by CIA and available commercially, the book is now a staple in many analytic training courses.

---

b. See Malcolm Gladwell, *Blink: The Power of Thinking Without Thinking* (New York: Little Brown and Company, 2005), 3-8. More recently, Gladwell told a story with a different outcome: In 2007 and 2008 the head of the investment firm Bear Stearns could not recognize, even in retrospect, the role his skills—the very skills that lay behind his earlier successes—played in his company's collapse. See Gladwell, “Cocksure: Banks, Battles, and the Psychology of Overconfidence,” *New Yorker*, 27 July 2009.

---

## Acknowledgements

I would be remiss if I failed to thank those who made this odd project possible, both now and 25 years ago: Mary-Ann Rozbicki, Paul Corscadden (both of whom have since passed away), and Carla Scopelitis—the first time around; and Carmen Medina, until December 2009, the director of the Center for the Study of Intelligence (CSI), which originally published this monograph, and CSI's Andy Vaart for its return. I would be doubly remiss if I failed to mention my gratitude to Jack Davis, the Directorate of Intelligence's long-time stealth subversive. Jack is one of the few people inside the Central Intelligence Agency who has sought to develop a genuine doctrine of intelligence analysis. His roller-coaster career—for a while in the 1970s and 1980s, it seemed that Jack was about to be fired at one moment, and at the next he was getting a medal for outstanding work—is in part a reflection of how difficult such a quest can be.

For me, Jack's early work with electronic mail (years before any of us had heard of the Internet) sensitized me to the medium's possibilities, including its culture-changing potential. He had kind words about my monograph from the outset, and he and I have had countless fruitful exchanges over the years. Finally, it was Jack who put the idea of reissuing the monograph into my head. He thus can claim a big share of the credit for the publication of this strange hybrid document; any shortcomings, of course, are my responsibility alone.

*tive securities cannot. They're always looking into an uncertain future, "anticipating what average opinion expects the average opinion to be," as Keynes put it.<sup>a</sup>*

And for its part, the Intelligence Community has taken a long list of initiatives ranging from structured analytic techniques to on-line blogs and social networking sites.

So why reissue this monograph? What could anyone gain from a 25-year-old piece on a subject that gets such broad coverage in today's popular literature? For me, the most telling answer came from a couple of talks based on the monograph that I gave in the spring of 2009. I spoke to conference-room-sized groups of analysts from CIA's Directorate of Intelligence (DI), and each time I had the impression that the talk was a revelation for most of the audience. Even after allowing as best I could for my own bias, I came away wondering whether information about cognitive science had percolated as widely as I had thought. In addition, in my regular interactions as a tutor of analysts and managers of analysts, I have similarly sensed that we have not absorbed the science into the way we think about our analytic jobs. Thus, one further attempt at consciousness-raising might not be out of place, especially since we have so many more ways to present this paper than we did 25 years ago.<sup>b</sup>

Moreover, while acknowledging that we have learned a great deal since 1984, I would argue that the elements of cognitive science highlighted in the monograph are still the ones of first-order relevance for the DI. I do not think an intelligence analyst will gain much professionally from knowing how neurons fire or which parts of the brain participate in which mental operations. I do consider it essential, however, that we be aware of how our brains ration what they make available to our conscious minds as they cope with the fact that our "ability to deal with knowledge is hugely exceeded by the potential

knowledge contained in man's environment."<sup>c</sup> Not only do they select among outside stimuli, they also edit what they let us know about their own activities.<sup>d</sup> This is the focus of the monograph.

### A Quick Summary

The monograph has two parts: first, a survey of cognitive science as we understood it in 1984; second, suggestions for changing the way we do intelligence analysis in light of what the discipline was telling us. As I have indicated, I think the survey section holds up pretty well. While I would like to think the reader will learn something useful from immersion in all the detail (notably the diagram on page 10, which makes graphic the many elements that interactively shape our conscious mental activity), the basic concept is quite simple. The conscious mind cannot track more than about seven cognitive elements at the same time (cognitive-science jargon often refers to these elements as *chunks*); and to cope with this constraint, our brains constantly manipulate those elements, always at top speed and usually outside our conscious awareness.

The second section is more of a mixed bag, but I believe the following concepts remain relevant:

- The importance of bringing to light what might be called, with a bow to former Defense Secretary Rumsfeld, the "unknown knowns"—the factors in our analysis that we are unaware of.

---

c. Jerome Bruner, *On Knowing: Essays for the Left Hand* (Cambridge, MA: Harvard University Press, 1962), 240.

d. A recent book by two neuroscientists puts it this way:

*Your brain lies to you a lot. . . . [It] doesn't intend to lie to you, of course. For the most part, it's doing a great job, working hard to help you survive and accomplish your goals in a complicated world. But because you often have to react quickly to emergencies and opportunities alike, your brain usually aims to get a half-assed answer in a hurry rather than a perfect answer that takes a while to figure out. Combined with the world's complexity, this means that your brain has to take shortcuts and make a lot of assumptions. Your brain's lies are in your best interest—most of the time—but they also lead to predictable mistakes.*

See Sandra Aamodt and Sam Wang, *Welcome to Your Brain* (New York: Bloomsbury, 2008), 2.

---

a. Justin Fox, review of two books about Keynes, *New York Times Book Review*, 1 November 2009, 13.

b. In 1984 CSI had no Web presence internally or externally, and the only way of delivering the findings of its fellows was through printed products advertised in internal notices or generated from searches conducted by library researchers.



- The conservative bias of our default analytic approach: we tend to work from what we already “know”—even though, in fact, a host of “unknown knowns” affect the “known.”
- The solitary nature of the writing process and the difference between the way intelligence analysts typically do their work (linear, cerebral, mostly written) and the way policymakers do theirs (nonlinear, transactional, mostly oral and interactive).
- The importance of constructing our prose with the reader in mind; the monograph, summarizing the work of cognitive-science investigators, states, “A striking aspect of the approach of skilled writers is the frequency with which they think about how they are affecting the reader.”
- The importance of the work done at the beginning of an analytic project—what the monograph calls the “conceptual front end.”
- The time needed to gain real skill at a craft; the monograph cites data suggesting that people are unlikely to get good at what they do for at least a decade.
- The use that can be made of information technology to improve the quality of analysis.

### **Suggestion: Develop Collaborative Analysis**

The Intelligence Community has not been blind to the potential in any of these areas. One has only to think of the emphasis on structured analysis and customer relevance; the attention we devote to scope notes; and the many blogs, the broader platforms such as A-Space and Intellipedia, and the coordination tools such as CIA’s POINT.

Since I wrote the monograph, however, I have nurtured an idea that I think would move our analysis well beyond these accomplishments. The monograph is not terribly clear on this point since I still was laboring to articulate the idea; but as far back as 1984, I was sure that electronic interaction was the wave of the future for DI analysis. More specifically, I thought the time was at hand when

we would be producing finished intelligence not just online but collaboratively.

The vision was reinforced later on when I read the introduction to *Thinking in Time*, by Harvard professors Richard Neustadt and Ernest May. The authors say they collaborated so closely that sometimes they lost track of who had written which part of the text.

*In every sense of the word the book is coauthored. We taught together, class by class, and have written together, chapter by chapter. It has been a long process. We each wrote half the chapters in our first draft, then swapped them for redrafting with the rule that anything could be changed. Almost everything was. We had a further rule that any changes could be argued. Some were, some not. We carried on the argument by means of successive drafts, with new material subject to the same rules. There were so many swaps that each of us would change something only to be surprised by the other’s “OK with me, you wrote it.” We debated sources and argued out interpretations. Not only can we no longer remember who first wrote what, we cannot now remember who first thought what—or even who first found what when we jointly researched something. We are of one mind and (we hope) one voice.”<sup>a</sup>*

I kept wondering, wouldn’t this sort of approach benefit our work?

I am sure that some collaboration of this sort is occurring today, but I have seen little sign, a quarter-century after the monograph came out, of any impact on the way the community as a whole operates. I have given a good deal of thought to the reasons my vision was so wrong:

- First, the idea required orders of magnitude more bandwidth and much more sophisticated software than we had in 1984. This constraint no longer holds, of course. We have plenty of

---

a. Richard E. Neustadt and Ernest R. May, *Thinking in Time: The Uses of History for Decision Makers* (New York: The Free Press, 1986), xii.

bandwidth, and programs to aid collaborative work are now widely available.

- Second, as I noted, the act of writing remains in essence a private process. To most writers, the idea of working interactively will seem like an intrusion into a space where they have always been alone.
- Third, the proposition faces serious cultural and organizational barriers. Hierarchy is at the core of any bureaucracy, including the DI's. The core of an analyst's performance evaluation is his or her publication record as an individual, and the serial review process is firmly ensconced as the way we generate our analytic product. A collaborative process would challenge the traditional approach in both areas. Managing it would require a healthy dollop of faith that something so messy and free-form can produce worthwhile results.<sup>a</sup> Furthermore, the obstacles loom larger in government bureaucracies than in the private sector. The risk/reward relationship is clearer in the latter, the budgetary process is more flexible and more under the control of management, and market competition and the bottom line will often prod managers toward innovation.<sup>b</sup>
- Fourth, interactive forums, whether Facebook and YouTube outside the community or Intel-

---

a. In the late 1970s the DI tried something analogous to this model (minus the IT element) in one of the offices responsible for political analysis. The branch-chief level of supervision was abolished, and analysts were encouraged to work on projects of their own choosing. A periodical called *Contra* was established to air alternative views. But over the next three or four years a sort of bureaucratic regression to the mean took place, and hierarchy and the review process reestablished themselves. *Contra* withered on the vine. The pendulum had swung far in the other direction by 1982, when Robert Gates (now the Secretary of Defense) took charge of the directorate. Gates undertook to review every draft himself and he pretty much did so.

b. The following is taken from Marc Ambinder, "Shutdown of Intelligence Community E-mail Network Sparks E-Rebellion," Blog: TheAtlantic.com, 6 October 2009:

*A former chief technology officer at the Defense Intelligence Agency . . . [stated,] "in some cases we are seeing IT departments cancel everything associated with innovation—which would be a sign of a dying organization in the private sector."*

lipedia and A-Space within it, may be enlightening for the participants, but nothing about them presses participants toward consensus or closure. Their mode is conversational; their strength lies in information-sharing and an ongoing batting-around of ideas. A participant might take insights reached online and use them elsewhere, but the idea of a discrete *product* coming out of such mediums seems almost a contradiction in terms.<sup>c</sup> The incentive to generate such a product would have to originate somewhere else.

- Finally, managers may have special difficulty adjusting to the interactive world. Lee Rainie of the Pew Internet and American Life Project distinguishes "nine tribes of the Internet." These, he says, can be divided into two groups: those "motivated by mobility" and the "stationary media majority." One component of the latter group is the "desktop veterans," who by themselves comprise 13 percent of the user population. Desktop veterans have been Internet users for 10.5 years on average; they are heavy users at home and work; 77 percent have cell phones. They are "content to use a high-speed connection and a desktop computer . . . [and] happy to be connected while they are stationary and sitting. So, they place their cell phone and mobile connections in the background. And their 2004 cell phone still serves its primary purpose for them—making phone calls. Online hit its zenith about 3-5 years ago when they first got broadband connections."<sup>d</sup> I myself am a charter member of this cohort.

I would guess that the older and more senior the manager, the more closely he or she will fit the "desktop veteran" profile. Such a manager will use electronic mail, word-processing software, spreadsheets, and presentational tools like PowerPoint. But he or she will have only second-hand familiarity

---

c. Whether a policymaker might find the online interchange useful is another question, one that is not addressed here but might be worth experimenting with.

d. Lee Ramie, "The Nine Tribes of the Internet." Presentation at the Washington Web Managers Roundtable, 10 June 2009.

with the interactive resources underlying any effort at online collaboration. (These will be more familiar to those in the lower ranks who are “motivated by mobility.”) And if it is correct, as noted above, that (1) interactive media exchanges are unlikely to reach closure without outside encouragement, and (2) managers, the obvious source of such encouragement, are not familiar enough with interactive media to provide it, the odds are against the full exploitation of this resource.

Perhaps we need not worry about this. Or perhaps the requisite organizational and (even more important) cultural changes seem too costly and the benefits too uncertain. I would only point out that cognitive-science literature makes clear the shortcomings of the process now in use, which amounts to an end-to-end series of solo efforts to get a piece drafted and then to coordinate and review it.

For every analyst and every reviewer in this serial process, the analysis starts from a body of analogies and heuristics that are unique to that individual and grow out of his or her past experience—after-images of ideas and events that resonate when we examine a current problem, practical rules of thumb that have proven useful over time.

The power of this approach is incontestable, but we are all too easily blinded to its weaknesses. The evidence is clear: analysis is likely to improve when we look beyond what is going on in our own heads—when we use any of several techniques designed to make explicit the underlying structure of our argument and when we encourage others to challenge our analogies and heuristics with their own. Little about the current process fosters such

activities, it seems to me; they would be almost unavoidable in a collaborative environment.

### ***Suggestion: Increase Cognitive Diversity***

One final thought, which is based on little more than observations about myself and those around me: I call it “cognitive diversity.” I believe the DI has always been populated very largely by serial thinkers like me, who analyze a problem by deconstructing it and laying out the result in writing, a quintessentially serial medium. I would of course be the last to decry this approach, but sometimes I am aware of getting so transfixed by my discoveries in the weeds that I have trouble getting back to the whole picture, much less the “so-what” that is the real purpose of the analysis. On the other hand, I have a good friend (he doesn’t work in intelligence) who as far as I can tell has not a serial bone in his body. Sometimes, however, he can see the entirety of an issue when I am still back in the weeds, and sometimes he runs circles around me when it comes to the “so-what.” In short, the analytic route taken by his mind differs from mine but is not necessarily less “analytic” in its own way, and sometimes he winds up in a more useful place than I do.

Might the DI do a better job if it were more cognitively diverse—if it took in more people with different cognitive “furniture”? I have opinions but no real answer to this question. I do believe diversifying the workforce in this way would require a cultural shift at least comparable to that involved in a shift to online substantive collaboration. Without such a shift, the directorate, like any organism under threat, would identify people who failed to fit the dominant pattern as foreign bodies and extrude them.



## *Thinking and Writing: Cognitive Science and Intelligence Analysis*

*By Robert S. Sinclair*

---

### **Summary**

The pages that follow describe some of the powerful metaphors about the workings of our minds that have developed over the past two decades, and attempt to apply those metaphors to the work of CIA's Directorate of Intelligence.<sup>a</sup> I believe the DI can thereby gain insights into such issues as the following:

- What is the best way to reconcile the bureaucratic imperatives of accountability, centralization, and structure with the fact that analytic work is essentially an individual effort?
- Can the directorate do a better job of pitching its analysis to catch the attention of its audience without sacrificing essential analytic detail?

I am not claiming that cognitive science offers definitive answers to such questions, but I do think it has something important to contribute to our understanding.

The term "cognitive science" embraces several disciplines, notably computer science, linguistics, and neurophysiology, as well as psychology. A cognitive scientist seeks to understand what the mind does when it searches for patterns, when it makes a value judgment, when it must choose between pattern-finding and judgment-making, when it engages in the myriad other activities that occupy it. Some fragmentary answers to questions such as this have become possible in the last 20 years.

Before they are six years old, nearly all humans learn to generalize, to impute continuity, to discern relationships, and to determine cause-and-effect. Moreover, we can store the conclusions drawn from such processes in a way that gives us access to them without burdening our working memory. We also learn a language, that uniquely human capacity which sits at the center of conscious cognitive activity. Language opens the way to abstraction and generalization, and permits each normal human to develop a rich network of concepts.

All of us are aware of the limitations of these processes. For example, we all are obtuse in dealing with logic and probability; we are comfortable with imprecision; and our minds are conservative in their approach to new information—quicker to recognize the familiar than the unfamiliar, reluctant to change concepts once we have accepted them. Finally, there are innumerable processes that influence our mental activity but are not accessible to the conscious part of the mind.

Some of the attributes that look like limitations, however, are actually the main sources of the mind's power; sloppiness is not just the bane but the strength of our mental activity. This is because of the role played by heuristics. The heuristic approach is a form of intelligent trial and error, in which we use experience and inference to refine a problem and render it workable. Few would give the process high marks for elegance, but it is quick, it gets the job done, and it keeps us from getting paralyzed by the range of choices confronting us. There are many disadvantages: the approach is inherently conservative, it tends to be imprecise, and it is not particularly congenial to logic, probability, and the scientific method. Nonetheless, heuristics are likely to remain the way we go about our busi-

---

a. The original monograph and this edition's introduction focus on CIA's Directorate of Intelligence because it is the analytic component of the Intelligence Community I am most familiar with. I hope, however, that the points I raise have relevance elsewhere in the IC's analytic world.

ness most of the time, particularly when we are deeply engaged in an issue.

The heuristic approach is based in part on deeply set mental patterns. “Working memory,” the part of the mind that does our conscious mental work, can handle about seven items at a time. In compensation, it can manipulate those items with extraordinary speed. Cognitive scientists refer to this manipulative capability as the mind’s *chunking* capacity—our ability to develop conceptual entities or *chunks*, to build hierarchies of those entities, to alter them, and to bring wildly differing entities together. We form chunks about any information that interests us, and we tend to believe our chunks are valid until the evidence that they are not is overwhelming. Each new bit of data is evaluated in light of the chunks already on hand; it is much harder to evaluate existing chunks on the basis of new evidence.

When we need to get through large quantities of data, when we do not have to move too far from an experiential reference point, and when a “best possible” solution suffices, heuristics and chunking can be amazingly effective, as Herbert Simon proved in his studies of first-class chess players. Such players are distinguished by the large number of board patterns (50,000, say) they keep in their long-term memories. Talent obviously is important as well, but Simon concluded that no one can become an expert player without such a store of chunks. Developing such a store in any field of mental activity is laborious, and there apparently are no shortcuts: the investment may not pay off for a decade.

All this information quickly takes on operational significance for the Directorate of Intelligence when we turn to writing, an activity that is simultaneously at the heart of the DI’s work and at the frontier of cognitive science. When cognitive scientists refer to the means used by humans to communicate with one another, they tend to use the term “language,” and unless they state otherwise, the word means speech rather than writing; few of them have focused on writing as a subject of research. Yet there are many ways in which the cognitive processes involved in writing differ from those involved in speech. Among other things, writing is capable of

far more breadth and precision; neither complex ideas nor complex organizations would be possible without it. On the other hand, because everyone works with speech, whereas not everyone works very much with writing, speech is a far more general medium of exchange. I would argue that speech is the medium to which all of us, even the compulsive writers among us, turn when an issue engages our emotions; and above all, I think, speech is the medium of decisionmaking.

I find it impossible to avoid the conclusion that our work will do its job better if it includes an element of speech—if we aim for prose that has a conversational ring. Such prose would often differ from that now produced in the directorate, and I think it would strain the organization to turn in this direction. I do believe, however, that to do so would help us get our message across.

Conversational prose must be produced by the original writer; it cannot be edited in later. This implies there will be an additional burden on the writer, but it is easy to overstate the burden and misstate the way it would be felt. Cognitive science makes it clear that, although the writing process is extraordinarily convoluted, good writers represent the writing problem as a “complex speech act.” With them, the conversational element is already present to some degree.

In the words of one expert, “a writer in the act is a thinker on a full-time cognitive overload.” One principal source of the overload is that the writer creates a datum—a malleable entity outside the mind that grows out of the mind’s internal workings. The datum acts as an extension of working memory, but working memory itself cannot keep track of all aspects of the datum unless it shifts constantly from one aspect to another. Skilled writers have various ways of reducing the overload. For example, like the chess master with his 50,000 patterns, they use chunks stored in long-term memory. Or they satisfice—an ungainly bit of jargon referring to the mind’s ability to accept a “best-possible” solution, at least temporarily. One of the most effective techniques is to develop and monitor a variety of heuristic strategies. Typically there are three sorts of strategies:

- A scheme *To Do*, the overall rhetorical problem as posed by the person solving it.
- Subordinate to this scheme, a scheme *To Say*—the substantive points to be made.
- Finally, in coordination with the scheme *To Say*, a scheme *To Compose*, the interaction between ideas and the developing text.

Good writers, it has been found, spend much of their time considering schemes *To Do* and *To Compose*; unskilled writers concentrate on schemes *To Say*. A striking aspect of the skilled writer's approach is the frequency with which he or she thinks about the audience. The cognitive-science literature indicates that developing a vivid image of the audience tends to enhance substantive content as well. Thus, being clear about the overall strategy and refining one's concept of the audience offer a way to bridge the gap between speech and writing.

Cognitive science also sheds light on the tension between creativity and the demands of a structured organization like the DI. The problems can be re-

duced if we recognize the overriding importance of what we do right at the outset of the analytic process. It is here that the writer makes assumptions about the overall rhetorical problem—the strategy *To Do*. I suspect this strategy is not well articulated for much of the DI's work, and I think it is problems at this level that cause the real headaches. There is, moreover, a built-in potential for conflict between what happens at the “conceptual front end” and the demands of a necessarily serial review process: plans *To Do* may have to be articulated over and over as the process goes forward.

Plans *To Do* typically get articulated incrementally and heuristically in any case, and thus it takes time to find out what is going on. Not all analysts find such inquiries congenial, and managers at all levels begrudge the time. But to the extent that all concerned can work early and often at narrowing the gaps between various concepts of the scheme *To Do*, the chance of a bumpy review process will diminish, whatever creativity there is may be preserved, and a sense of collegiality will be enhanced.

### Recommendations

These findings have obvious implications for the way the directorate recruits and trains its people. They suggest that more emphasis should be given to effectiveness at the *To-Do* level, and perhaps that skills of the *To-Say* variety are somewhat less important. Moreover, if it is true that around 10 years are needed to acquire a first-class network of chunks, training probably should figure far more heavily in the DI's thinking than it does now.

The trends already under way in the DI should encourage conversational prose, and the notion of bringing the audience into focus should help writers and reviewers establish common ground to work out the nonsubstantive aspects of a paper. It still will take repeated effort, however, to bridge the conceptual gaps that are bound to appear.

The conflict between organizational imperatives and the way analysts do their cognitive work might

be reduced by divorcing substance from hierarchy as often as possible; through hierarchy-jumping in contacts with consumers; and by enfolding SAFE—which has the potential to bring significant changes in the way the directorate does its work—in the culture of the directorate. Concept papers could be made to bring out the *To Do* schemes of the various participants more clearly.

The directorate might explore the possibilities of non-written media, such as television, more extensively.

Finally, the ideas contained in this essay do not exhaust the possibilities offered by cognitive science. Suggestions for further work include a cognitive task analysis of the analytic process and an exploration of the extent to which the directorate's concentration on the written word limits its analytic flexibility.

## Preface

This essay has its origins in the controversies over CIA's analytic work that arose after Vietnam and Watergate, the criticisms levied against the Agency by its principal consumers, and the Agency's efforts to respond to those criticisms. For much of this period it seemed that producer and consumer were talking past each other, and I sometimes thought that the resultant organizational tinkering was proceeding with little reference to the analytic work going on around me. The producer-consumer gap has since been greatly narrowed, and our analytic work seems more on target now than in the past. I have a sense, however, that we still do not understand the way we do our analysis, or the intricacies of the producer-consumer relationship, as well as we might. These are the questions

addressed in the following pages. The essay is a layman's view of a body of knowledge that did not exist two decades ago, together with an attempt to lay out some operational implications of that knowledge.

The concerns that stimulated my inquiries still were pretty inchoate when I submitted my request for a sabbatical (in this sense the genesis of the essay exemplifies the cognitive processes it discusses). I think it is a tribute to those who approved the request—perhaps to their faith, perhaps to their gullibility—that I have been allowed to spend six months finding out what my questions were and then trying to answer them.

## Introduction

Two quotations sum up what this essay is about:

*Our insights into mental functioning are too often fashioned from observations of the sick and the handicapped. It is difficult to catch and record, let alone understand, the swift flight of a mind operating at its best.<sup>a</sup>*

*A writer in the act is a thinker on a full-time cognitive overload.<sup>b</sup>*

The pages that follow will be concerned with the “mind operating at its best”—and in the Directorate of Intelligence, which operates mainly in the written mode, that usually means a mind “on a full-time cognitive overload.” In brief, I hope to describe some of the powerful metaphors about the workings of our minds that have developed over the past couple of decades. I think these metaphors can help provide better answers to such questions as the following:

- What is the best way to reconcile the bureaucratic imperatives of accountability, centralization, and structure with the fact that analytic work is essentially an individual effort?
- Can the directorate do a better job of pitching its analysis to catch the attention of its audience without sacrificing essential analytic detail?

The hardest part of the essay to get right comes at the outset. Few readers are likely to know much about the territory I am asking them to explore, and much of the terminology—“cognitive science,” “psycholinguistics,” “neurophysiology”—is abstruse and off-putting. I believe, however, that the new territory is likely to prove surprisingly accessible. One of the virtues of cognitive science, in fact, is the way it gives insights into what we have been doing all along. Unlike Moliere's *bourgeois gentilhomme*, who was surprised to learn he had been speaking prose all his life, I think most of us will recognize ourselves in what follows.

But how relevant are these findings to the Directorate of Intelligence? Many might argue, for instance, that introspection of this sort will gain nothing.

---

a. Bruner, *On Knowing*, 15.

b. Linda Flower and John R. Hayes, “The Dynamics of Composing: Making Plans and Juggling Constraints” (article provided by the authors).



ing at best and paralysis at worst. The point is just to do the analysis, and doing it is learned on the job, in consultation with our consumers. Those who hold this view probably would applaud the thought behind the following bit of doggerel:

*A centipede was happy quite  
Until a frog in fun  
said, "Pray, which leg comes after which?"  
This raised its mind to such a pitch  
It lay distracted in a ditch  
Considering how to run.*

This is not an easy argument to dispose of, and the centipede will reappear from time to time in the pages that follow. The argument is based on the proposition, which seems correct to me, that the DI is essentially a craft shop, and I agree that most of the learning of a craft does take place on the job. But look again at the quotation from Flower and Hayes that I began with. If the very act of writing puts a writer—any writer at all—into “full-time cognitive overload,” then perhaps we would benefit from a better understanding of what contributes to the overload.

The novelist and poet Walker Percy offers a concept that may be even more fruitful. In a series of essays dealing with human communication, Percy asserts that a radical distinction must be made between what he calls “knowledge” and what he calls “news.”<sup>a</sup> Percy’s notion takes on added significance in light of the findings of cognitive science (of which he seems largely unaware), and I will be discussing it at greater length in due course. For the present I would simply assert that the nature of our work forces us to swing constantly back and forth between knowledge and news, and I believe cognitive science has something to contribute to our understanding of the problem. I am not claiming it offers

a panacea; I do think that in ways such as this it sheds powerful light on important practical issues.

A few clarifications are in order at the outset. First, although this essay talks a lot about writing, it is not designed to deal with the how-to-write issue. As the title indicates, its topic is thinking and writing—the complex mental patterns out of which writing comes, their strengths and limitations, and the challenges they create, not just for writers but for managers. I hope my suggestions are relevant to the never-ending struggle toward better writing, but I am trying to cast my net more broadly than that.

Second, it should be obvious that cognitive science is only one of many ways to approach human mental activity. Factors other than the cognitive activities discussed in this essay play major roles in the way we do our mental work—for example, the attitudinal predilections that can be measured by psychological testing, or the emotional factors discussed in the psychoanalytic literature. To my mind, however, the findings of cognitive science have special relevance for the way the directorate does its business because they illuminate the process itself as well as the equipment we bring to the procedure.

Third, I am not claiming to speak as an across-the-board expert in cognitive science. What follows is the result of very rapid chunking (a useful bit of cognitive-science jargon, as we shall see), and it tends to emphasize the work of Herbert Simon and his colleagues at Carnegie-Mellon University. I think there is good reason for such an emphasis since I believe their work has particular relevance for the DI’s writing-based culture. I would be the last, however, to assert that what follows is a definitive treatment. That is why many of the suggestions in the last section are put as questions rather than assertions. My main goal is to start a discussion; any progress will need the help of real experts.

---

a. Walker Percy, *The Message in the Bottle* (New York; Farrar, Straus, and Giroux, 1982).

## *I. A Quick Survey*

First we need to get some notion of what is involved in cognitive science. The term embraces several disciplines, notably computer science, linguistics, and neurophysiology, as well as psychology. Very broadly there are three main areas of inquiry: how the various parts of the brain (and the nervous system as a whole) interact, in both a neurological and a functional sense; how the human capacity for language is turned into specific linguistic skill; and how people analyze and solve problems, both simple and complex. To pick a starting point, we might say the field began to open up in 1959, when the linguist, and all-round radical thinker, Noam Chomsky published a scathing review of a book by B.F. Skinner, whose behavioral-science approach had dominated research into the workings of the mind up to that point.

### *Viewing the Interstate*

To get some flavor of what this activity has produced, imagine that you are sharing my aerie on the 10th floor of a building that overlooks Interstate 66, just outside Washington, DC. Every afternoon we can watch a game involving the Virginia State Police and that rugged individualist, the Washington commuter. No car is supposed to use the westbound lanes of the interstate after 3:30 p.m. unless there are two or more people in it, but some commuters tempt fate every afternoon, and many afternoons the police are waiting for them as they come over the rise on the access ramp.

From our vantage point we can view this game in different ways. We can guess whether a specific car is going to be stopped. We can see if we can discern any patterns in the frequency with which the police appear. Or we can ponder the values that come into play when a government limits the use of a road and individuals decide to ignore the rules. In the latter case we can look at the conflict more or less dispassionately—guessing, perhaps, about the likelihood that the government will stick to its position—or we can join the fray, siding with either the government or the commuter. I would argue that all

these activities are analogous to the sort of work DI analysts do at one time or another.

Now imagine that there is a third person, a cognitive scientist, in the room with us. To this person the interesting question is not the judgments you and I make about what is happening on the interstate, but the processes our minds use to make those judgments. He or she wants to know what the mind does when it searches for patterns, when it makes a value judgment, when it is forced to choose between pattern-finding and judgment-making, when it engages in the myriad other activities that occupy it. And the amazing thing—amazing given the physical complexity of the brain, where the neocortex alone contains something like 10 billion nerve cells, each capable of firing several times a second—is that some fragmentary answers are possible.

To illuminate the answers to the scientist's questions, we need to operate at a more elemental level. As you imagine yourself looking at the interstate, consider not the little drama that is going to begin when the next car gets to the top of the rise, but your perception of the car itself. A crucial point is that perception itself involves analytic work of a very basic sort. You constantly check out the characteristics of the car and match this information with information already in your memory. You then predict what the car is going to do next, and you check that prediction against what happens. You are not aware of these processes because long ago you learned them thoroughly (cognitive scientists would use the term "overlearned") and no longer have to waste the limited capacity of your working memory on them.

But it is worth pausing to reflect on what we all learn without apparent difficulty before the age of six. In the context of the I-66 example, any of us would know that the car that just emerged from behind the building across the way is the same one that disappeared behind the other side a few moments ago; that the term "car" can be applied to both the Continental and the Honda, however different they may appear, but not to the GMC pickup;

that the shiny metal and glass, glimpsed momentarily through a gap in the fence, is another passing car; and that the styrofoam cup in the road is being blown about because a car has just passed. In more formal terms, we all learned to generalize, to impute continuity, to discern relationships, and to determine cause-and-effect. And we can store the conclusions drawn from such processes in a way that gives us access to them without burdening our working memory. This is a formidable array of skills (they obviously are the basis for the skills used in the DI), and they all developed with such ease that the childhood process seems almost automatic.

But this is not all. We also learned to give names both to the things we see and to the relationships among those things. In short, we learned a language, the uniquely human capacity that sits at the center of all our conscious cognitive activity. Obviously it is language that permits us to communicate information of any complexity to each other, but in addition, language is a prerequisite for the steadily increasing complexity of which our own minds are capable as they develop. Naming a thing and naming its relationship with other things are themselves acts of abstraction and generalization, and once we have taken this step we can name relationships among relationships, thereby building a dense conceptual network.

All of us know, however, that if we did nothing but marvel at the achievements of our minds we would be leaving a lot unsaid. I am not referring to the times our mental processes prove inadequate to a task; I want to focus on the limitations these processes encounter even when they are working well—limitations, we often see more easily in others than in ourselves. For example, we all are determinedly obtuse when we attempt to deal with logic and probability. We are much readier to use them to buttress arguments we have already worked out than to discipline those arguments.

We also are surprisingly comfortable with imprecision. Douglas Hofstadter, in his brilliant and infuriatingly self-indulgent book, *Gödel, Escher, Bach*, states the situation well:

*The amazing thing about language is how imprecisely we use it and still manage to get away with it. SHRDLU [an artificial-intelligence computer program] uses words in a “metallic” way, while people use them in a “spongy” or “rubbery” or even “Nutty-Puttyish” way. If words were nuts and bolts, people could make any bolt fit into any nut; they’d just squish the one into the other, as in some surrealistic painting where everything goes soft. Language, in human hands, becomes almost like a fluid, despite the coarse grain of its components.<sup>a</sup>*

A small example from close to home: a group of DI analysts, asked what numerical odds they would associate with the word “probable,” gave answers ranging from 50 to 95 percent. Part of the tendency toward imprecision derives from the way humans surround every explicit statement with a cloud of assumptions. Our cognitive activity would be cumbersome indeed if we had to articulate all these assumptions, even though in not doing so we may discover afterwards that we were not talking about the same thing as someone else. Another factor making for imprecision is the element of abstraction that is built into language, since with abstractions it is difficult to be clear exactly what we are referring to. Indeed people deal in hierarchies of abstractions, and the DI generally operates toward the more abstract rather than the more concrete end of the hierarchies.

Moreover, our minds are conservative in the way they select information for processing by working memory. The capacity of working memory is tiny: It can only deal with about seven items of information at once. Yet the mind must sort through a welter of sensory data (held in very short-term buffers and all clamoring for attention) and must also make effective use of the information in long-term memory. Faced with this cacophony, our minds generally are quicker to recognize the familiar than the unfamiliar, and data already in our memories heavily influence the processing of new data. We rarely are fully conscious of the choices that are made. To put it more

---

a. Douglas Hofstadter, *Gödel, Escher, Bach: An Eternal Golden Braid* (New York: Vintage Books, 1980), 674–75.

concisely, within broad limits, we see what we want to see and make what we want of what we see.

Complex though they are, moreover, these aspects of the mind's activities, which are at least partly conscious and volitional, are far from the whole story. The literature is full of information on processes of which we cannot be aware, processes that constantly affect our conscious mental operations.<sup>a</sup>

---

a. These processes have been discovered through neurological investigations; they are distinct from the unconscious activities investigated by psychoanalysts. For further information, see Erich Harth, *Windows on the Mind* (New York: Morrow, 1982) and Gazzaniga and Ledoux, *The Integrated Mind* (New York: Plenum, 1978). Among other things, the latter book pro-

But given all this uncertainty, all these aspects of our minds that are in principle unknowable or beyond our control, what can introspection gain us? Don't you wind up in the ditch with the centipede if you start asking what goes on in our minds? In fact, however, it is by looking at some of these very attributes that we gain useful insights into the way we go about our work. Some things that up to now have looked like weaknesses are actually the main sources of the mind's power. It turns out that sloppiness is not just the bane but the strength of our mental activity.

---

vides a useful corrective to the popular literature's treatment of the right brain/left brain phenomenon.

## II. Heuristics and the Incredible Chunk

Some readers may remember the following episode from Jacques Tati's classic film comedy, *Mr. Hulot's Holiday*. The scene is a French railroad station. The camera takes in three platforms, and a little knot of travelers is waiting expectantly on the platform in the center. Suddenly the loudspeaker overhead begins to bleat a long and utterly unintelligible announcement; and just then a train slowly approaches the platform on the left. After a few moments of growing agitation the travelers, goaded by more bleats from the loudspeaker, disappear down the platform steps. They emerge on the left platform just as it becomes obvious that the train is going through without stopping. *Their* train is at that moment pulling in at the platform to the right.

This episode can serve as a model for the way our minds work in a much larger sphere. It is not hard to reconstruct the cognitive work being done—done at a furious pace—by the travelers. First, with the trains, they are following the over-learned patterns that we followed with the cars on the interstate: matching what their senses tell them with information already in their memories, they recognize the train and make some predictions for their senses to check out. But they then apply a

similar process of perception/prediction/verification to information that for the most part is generated internally. What emerges is a collection of unspoken hypotheses, which might perhaps be put into words as follows: "My train is the next one scheduled; the loudspeaker almost certainly is announcing my train; that train over there seems to be the only one coming in; all trains must stop at the station; so the train over there must be mine." They easily reach a consensus; perhaps they even perceive that the train is slowing a bit (it's easy to let one's presuppositions affect one's perceptions). And, of course, they wind up looking silly.

*The  
heuristic approach is a  
form of intelligent trial and error  
in which we use experience and  
inference...*

But—and this is a critical part of the model—they do *not* miss their train. Further, I think anyone watching the movie would assume that even if the travelers make some mistakes, they still can catch the train if they scramble. In other words we and the travelers start from the idea that the system is built to accommodate a certain amount of trial and error.<sup>a</sup>

---

a. In New York City that is the case, according to the *New York Times*, which on 16 October 2009 reported that commuter trains almost always leave their platforms a minute after their scheduled departure times.

The jargon for the approach of the travelers is *heuristics*. A cognitive scientist might call their struggles a heuristic exploration of a poorly defined problem space; he or she would argue that heuristics offer a more satisfactory account of the way humans learn than the stimulus-response approach of the behavioral scientists. The heuristic approach is a form of intelligent trial-and-error, in which we use experience and inference to clarify, narrow, or otherwise refine a problem to make it workable. Logic has a role, but a subordinate one. The essence of the process is a non-random bargaining around. If we decide one tactic is not working, we back up and try another, but we are reluctant to do so. We tend to assume our theories are right until they are firmly disproven. Much of the time we end our search for solutions before we achieve the optimal outcome. We accept a “best-possible” solution and move on to something else, or to use the ungainly word invented (I think) by Herbert Simon, we *satisfice*.<sup>a</sup>

Few would give the process high marks for elegance, but it does deserve high marks for success. In other words, it works. Heuristics seem to have played a major role in the rapidity of human evolution, and they are equally prominent in the development of each individual. This is true of both cognitive and physical skills. The way children learn their native language is a particularly clear example. Preschool children do not learn a language by memorizing rules or parroting the sentences they hear. The essence of their approach from the very beginning is to experiment with linguistic patterns of their own. Using this approach they can learn rules of extraordinary complexity with little apparent trouble. Moreover, from the beginning, their linguistic activity has both a creative component and an element of satisficing. Logic and the scientific method—what Piaget called formal procedures—come late, are almost never fully assimilated, and seem to be used less systematically by nearly everyone after the age of 20 or so.<sup>b</sup>

---

a. The last syllable of the word rhymes with “dice.” Herbert Simon, a Nobel laureate in economics, began delving into cognitive theory partly out of dissatisfaction with the benefits-maximizing “rational man” of classical economic theory. His “thinking man” is a satisficer. See Simon, *Models of Thought* (New Haven: Yale University Press, 1979), 3.

b. See Carol Gilligan and J.M. Murphy, “Development from

What is it about heuristics that makes them so useful? First, they are quick and they get the job done, assuming the experiential base is sufficient and a certain amount of satisficing is not objectionable. Second, what cognitive scientists call the problem-space remains manageable. Theoretically that space becomes unmanageably large as soon as you start to generalize and explore: any event may be important now, any action on your part is possible, and you could get paralyzed by possibilities as the centipede did. But humans constantly narrow the problem-space on the basis of their own experience. And most of the time the results are acceptable: what more efficient way is there to narrow an indefinitely large problem-space?

But there are obvious limitations to the approach, and the limitations become more apparent when one faces issues like those confronting the DI:

- Heuristics are inherently conservative; they follow the tried-and-true method of building on what has already happened. When the approach is confronted with the oddball situation or when someone asks what is out there in the rest of the problem-space, heuristics begin to flounder. Yet we resist using other approaches, partly because we simply find them much less congenial, partly because the record allows plausible argument about their effectiveness when dealing with an indefinitely large set of possibilities.
- As most people use them, heuristics are imprecise and sloppy. Some of the reasons why cognitive activity is imprecise were noted earlier; another reason is the tendency to satisfice, which encourages us to go wherever experience dictates and stop when we have an adequate answer. With perseverance and sufficient information one can achieve considerable precision, but there is nothing in the heuristic approach itself that compels us to do so and little sign that humans have much of an urge to use it in this way. Most of the time,

---

“Adolescence to Adulthood: The Philosopher and the Dilemma of the Fact” in *Intellectual Development Beyond Childhood*, No. 5 in the series *New Directions for Child Development* (San Francisco, CA: Jossey-Bass, 1979).

moreover, the information is not terribly good. We then may find ourselves trying to get more precision out of the process than it can provide.

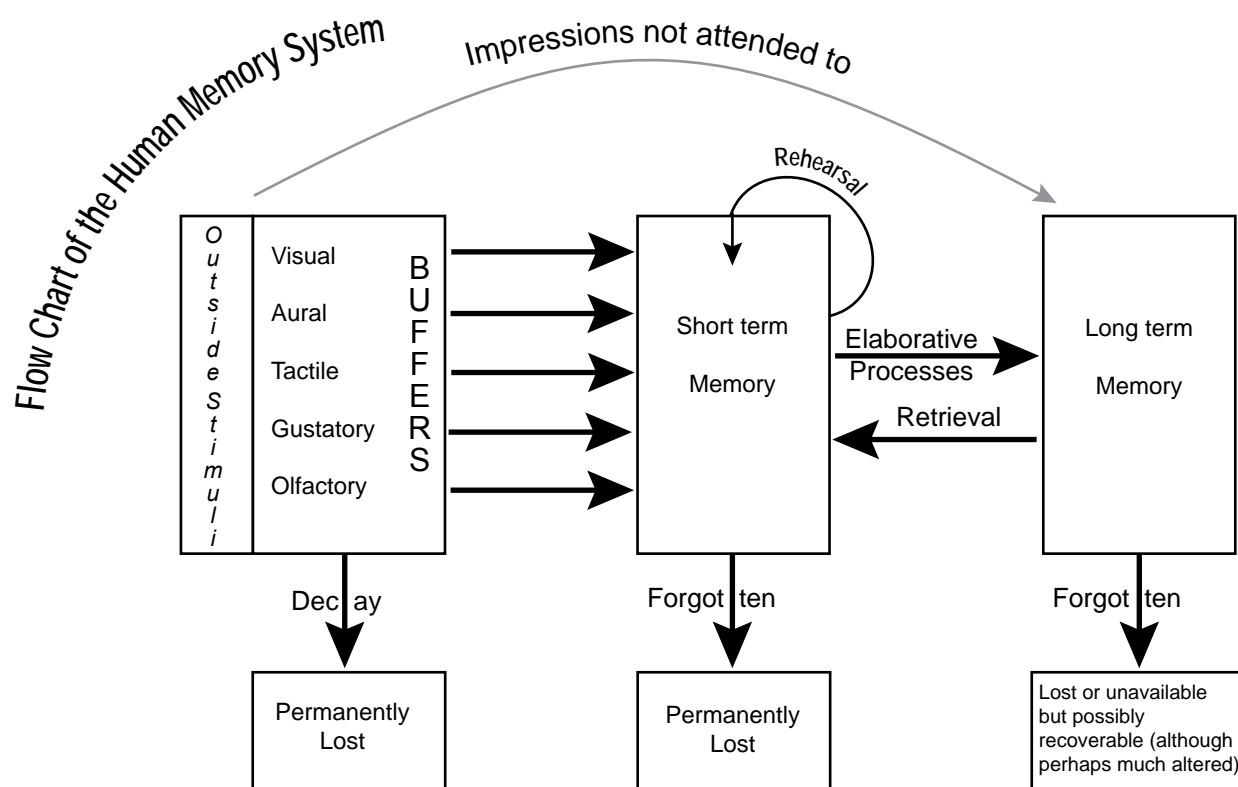
- In everyday use, heuristics are not congenial to formal procedures such as logic, probability, and the scientific method. This fact helps explain why we rarely use logic rigorously, why we tend to be more interested in confirming than in disconfirming a hypothesis, and why we are so poor at assessing odds.

Warts and all, however, heuristics are likely to remain the way we all go about our business most of the time. And it seems to me that the more deeply engaged we are, the likelier we are to operate this way. Perhaps it is an obvious proposition that humans use the approaches they are most comfortable with when an issue is important to them, and heuristics, whatever its drawbacks, are what we all are comfortable with.

One reason the heuristic approach is so deep-set is that it uses even more deeply set mental patterns. Once again at this level we find capabilities with great power (they let us run circles around computers in many respects) but with significant limitations as well—limitations we probably would be well advised to learn to work with rather than try to change.

The below diagram is typical of the way cognitive scientists represent the mind's operations.<sup>a</sup> We all have a system of buffers that enables us to organize incoming stimuli, a long-term memory of essentially infinite capacity, and a short-term or working memory that does our conscious mental work. At the level of generality at which we are operating, the roles of the buffers and long-term memory are fairly obvious. It is working memory that needs further discussion.

a. This particular diagram is taken from Morton Hunt, *The Universe Within* (New York: Simon and Schuster, 1982), 103.



The two salient features of working memory are its speed and its limited capacity. In one of P.G. Wodehouse's novels a particularly dunderheaded character is said to be capable of only "one thought at a time—if that," and the rest of us are closer to this standard than we might like to admit. As I noted earlier, working memory can handle perhaps seven items at a time. This number has to include not just substantive information but also any processing cues. If we have something we want to say in a complex sentence, for example, we must use one or more of the seven slots to keep track of where we are—which is one reason our spoken discourse rarely uses complex sentences.

It should be apparent the heuristic approach is critical to the effectiveness of our conscious mental activity, since short-term memory needs procedures like heuristics that narrow its field of view. On the other hand, the drawbacks are equally apparent. The ability to process large quantities of information is always an advantage and sometimes a necessity. How can we operate effectively if we can consider so little at a time? The answer to this question lies in the speed and flexibility with which we can manipulate the information in short-term memory; to use the terminology, in our *chunking* prowess.

In Morton Hunt's formulation, a *chunk* is

*any coherent group of items of information that we can remember as if it were a single item; a word is a chunk of letters, remembered as easily as a single letter (but carrying much more information); a well-known date—say, 1776—is remembered as if it were one digit; and even a sentence, if familiar ("A stitch in time saves nine.") is remembered almost as effortlessly as a much smaller unit of information.<sup>a</sup>*

A chunk, it should be clear, equates to one of the roughly seven entities that short-term memory can deal with at one time. Hunt's formulation notwithstanding, it need not be tied to words or discrete symbols. Any conceptual entity—from a single letter to the notion of Kant's categorical imperative—can

be a chunk. And not only do we work with chunks that come to us from the outside world, we create and remember chunks of our own. Anything in long-term memory probably has been put there by the chunking process. We build hierarchies of chunks, combining a group of them under a single conceptual heading (a new chunk), "filing" the subordinate ideas in long-term memory, and using the overall heading to gain access to them. We can manipulate any chunk or bring wildly differing chunks together, and we can do these things with great speed and flexibility.

Consider just two examples. First, the use you make of this essay. You are not likely to recall a single sentence of it when you finish; but you will (I hope) have stored ideas derived from it in your long-term memory. You will, in fact, have developed your own chunks on the basis of what you have read. If I have done my job there will be a great many such chunks, they will bear at least a familial resemblance to the ideas I am trying to explicate, you will have formulated them quickly and easily, and you will manipulate them on your own.

The second example is very different. Simon and his colleagues at Carnegie-Mellon University have specialized in the oral protocol, in which individuals are asked to articulate as much as possible of their thought-processes as they solve a problem or engage in some other kind of mental activity. Here is an excerpt from one such protocol, in which the subject has been asked to turn a simple story problem into an algebraic equation. The aspect that sheds light on the chunking process is the use of "it" at the two points where I have underlined the word.

*We'll call the number  $n$ . It says that if we multiply it by 6 and add 44 to it, "the result is 68."<sup>b</sup>*

You will notice that superficially the meaning of "it" changes between the first and the second use; the word refers to  $n$  the first time and to  $6n$  the second. But there is an underlying sense in which "it" remains constant. One might articulate this sense as "the idea (or chunk) I am currently manipulating."

---

a. Morton Hunt, *The Universe Within*, 88.

---

b. Simon, *Models of Thought*, 209.

ing, whatever its specific content at the moment.” Furthermore, this sense is clear enough for us to have little doubt what the person doing the protocol is trying to say. In this brief excerpt the subject of the experiment has been caught in the act of rapid chunking, and we have been caught in the act of comprehending the process without noticing it. The process was as quick for us as it was for the subject.<sup>a</sup>

In some ways “chunk” is a misleading term for the phenomenon. The word calls to mind something discrete and hard-edged, whereas the very essence of the phenomenon is the way we can give it new shapes and new characteristics, and the way conceptual fragments interpenetrate each other in long-term memory. A chunk might better be conceived of, metaphorically, as a pointer to information in long-term memory, and the information it retrieves as a cloud with a dense core and ill-defined edges. The mind can store an enormous number of such clouds, each overlapping many others.

This “cloudiness”—the way any one concept evokes a series of others—is a source of great efficiency in human communication; it is what lets us get the drift of a person’s remarks without having all the implications spelled out. But it can also be a source of confusion. Consider the following examples, some showing efficiency, some opening the way for confusion.

- My teenage daughter, nervously trying to pour conditioner into a narrow-necked bottle, says, “This isn’t going to work.” But she doesn’t hesitate in her pouring, and the job is completed without mishap. Her words, her inflections, and her actions have communicated a whole web of concepts to me (and perhaps to herself) which are far indeed from what she actually said.
- A boss comes out of an inner office, puts a small stack of typewritten pages on the secretary’s desk, says, “There aren’t enough copies,” and goes back into the office. The secre-

tary might tell you the boss had said to make more copies. The boss might agree. Both are working from assumptions that go unarticulated but not from what the boss actually said. (In addition you, the reader, may have thought of the boss as being male and the secretary as being female—another unarticulated assumption.)

- In a cognitive science lecture a speaker reads aloud the following sentences: “Mary heard the ice cream truck coming down the street. She remembered her birthday money and ran into the house.” Those in the audience conclude that the sentences are about a little girl who is going to buy some ice cream; the speaker notes that this conclusion is based on inference, not on anything explicitly stated.

Cognitive scientists apply a variety of terms—“networks,” “schemata,” “scripts”—to the unspoken contexts into which we set pieces of information. The last example in particular makes clear how eager we are to build such contexts. Neither we nor the lecturer’s audience had any obvious need to construct any story at all from those two sentences about Mary, yet nearly everyone does. We form chunks about any information that interests us, and (heuristics again) we tend to believe our chunks are valid until the evidence that they are not is overwhelming. And we form our chunks right away, as the information is coming in: there is no nonsense about waiting for a sufficiency of evidence. Each new bit of data is evaluated in light of the chunks already on hand, all of which are treated as at least arguably valid. We find it far more difficult to evaluate existing chunks on the basis of a new piece of data.

When we need to get through large quantities of data, when we do not have to move too far from an experiential reference point, when a “best possible” answer is sufficient (and these criteria are met most of the time), chunking can be amazingly effective. The approach works in contexts far more complicated than might be expected, as is indicated by the experiments with chess players conducted by Simon and others.

---

a. Another point this little excerpt brings to light is the degree of linguistic imprecision we can tolerate, particularly when we are dealing with spoken language



What are the hallmarks of a first-class chess player? Not, it seems, the ability to look further ahead in the game or to evaluate a larger number of possible moves; the master's approach is not much more extensive in these respects than the beginner's. What distinguishes the master and the grandmaster is the store of chess patterns built up in long-term memory over years of competition and study. Simon estimates that a first-class player will have 50,000 of these patterns to call on—by no means a small number, but orders of magnitude less than the theoretical possibilities that flow from any given position. The expert can use them to drastically reduce the number of choices he must consider at any point in a game, with the result that he often hits on an effective move with such speed that the observer attributes it to pure intuition.

So formidable is this ability (which is little more than heuristics and chunking) that humans were able to keep ahead of chess-playing computers for longer than most experts in the field predicted. When a computer that could beat nearly every human was developed, it did not duplicate the chunking approach but instead performed a brute-force search of all possible moves. Computers still are far inferior to humans in their ability to narrow the field heuristically, but they can now explore the whole field of chess moves in the time normally allowed between moves.

None of this emphasis on the importance of a store of chunks is to deny the importance of individual talent in chess or any other field of mental activity. There is nothing that says you will automatically become a grandmaster if you simply lodge 50,000 or more chess patterns in your long-term memory. But you apparently cannot become a grandmaster *without* such a store of chunks, and cognitive scientists argue that a network of similar size and complexity is essential for effective work in most fields. Developing an adequate network is laborious and time-consuming, and there apparently are no shortcuts. Edison's aphorism about genius being two percent inspiration and 98 percent perspiration is buttressed by J.R. Hayes's discovery that almost without exception, classical composers did not start

producing first-rate work until they were at least 10 years into their careers.<sup>a</sup>

By now we can speak in some detail about how the centipede runs—i.e., how humans do mental work. We still are speaking in highly schematic terms, the picture that emerges is exceedingly complex, and there is a great deal that cannot be addressed intelligently at all. But I think we can draw some illuminating conclusions, often humbling and heartening at the same time:

- Heuristics—non-random exploration that uses experience and inference to narrow the field of possibilities—loom large in the development of each individual and are deeply ingrained in all of us (particularly when we are doing something we consider important). Combined with the chunking speed of short-term memory, the heuristic approach is a powerful way to deal with large amounts of information and a poorly defined problem space.
- But there is always a tradeoff between range and precision. The more of the problem space you try to explore—and the “space,” being conceptual rather than truly spatial, can have any number of dimensions—the harder it is to achieve a useful degree of specificity. Talent and experience can often reduce the conflict between the need for range and the need for precision, but they cannot eliminate it. We almost always end up satisficing.
- We are compulsive, in our need to chunk, to put information into a context. The context we start with heavily conditions the way we receive a new piece of information. We chunk so rapidly that “the problem,” whatever it is, often

---

a. Mozart, who springs to mind as the obvious exception, turns out to have been a late bloomer by this standard. Hayes used a criterion for “first-rate work” that seems terribly unscientific at first: he decided that a composition qualified if five or more recordings of it were listed in the Schwann catalogue. The idea looks a bit more sensible at second glance, and in any case the consistency of Hayes's findings is striking. He obtained data for 76 composers and found only three—Satie, Shostakovich, and Paganini—who produced even one such piece within 10 years of the time they began studying music intensively.

has been sharply delimited by the time we begin manipulating it in working memory.

- Skill in a given area of mental activity involves both talent and time, the latter to build the extensive network of chunks that appears to be a prerequisite for really first-class work. The effort typically extends over several years; the investment may not pay off for a decade.
- Although the conceptual network formed through years of experience may make an individual a more skillful problem-solver, it can also make him or her less open to unusual ideas or information—a phenomenon sometimes termed “hardening of the categories.” The conservative bias of the heuristic approach—the tendency we all share of looking to past experience for guidance—makes it easy for an old hand to argue an anomaly out of the way. In fact the old hand is likely to be right nearly all the time; experience usually does work as a model. But what about the

situation when “nearly all the time” isn’t good enough? Morton Hunt recounts an instance of a computer proving better than the staff of a mental hospital at predicting which patients were going to attempt suicide.

- Finally, the mental processes outlined in this section seem certain to remain the processes of choice for all of us. With all their drawbacks, heuristics remain arguably the most efficient way to narrow a large field of possibilities; and efficient or not, we like it.

So far, however, the findings we have discussed are little more than abstractions; they lack an operational context for the DI. And for good reason: the directorate is basically a writing organization, and cognitive science can hardly be relevant to our work until it addresses the main thing we do. As the section that follows makes clear, the focus quickly becomes operational as soon as the ideas accumulated so far are applied to the issue of writing.

### III. Speaking and Writing, News and Knowledge

When people who deal with cognitive activity discuss the means we use to communicate with each other, they tend to refer to something they call “language.” Unless they specify otherwise they generally mean spoken language, not writing. They discuss written language much less frequently, and when they do there sometimes is even a hint that they consider it less important than speech. Few have focused on it as a subject of research. This seems strange, since not only do these people make their living by the written word, they would have little hope of fully articulating their thoughts unless they had writing available as a tool.

What accounts for this phenomenon? First, it simply may not have occurred to them that there is a useful distinction to be made between speech and writing. Second, writing is a much more opaque process than speech and thus is harder to analyze.

A speaker literally is making up what he says as he goes along, and it is relatively easy to gain insights into his cognitive activity by observing what he says and how he says it. The writing process is more convoluted and internal, and it is harder to get even a general notion of what is happening.

But I think there is a third reason, one that has great import for the way the directorate goes about

*Speech  
is the natural medium of  
decisionmaking...The essence of the  
decision process is oral.*

its cognitive work. In some ways those who treat writing as less important than speech are right. The spoken word is the dominant (usually the only) linguistic medium in those pre-

school years when our basic conceptual furniture is being established, and I would argue that it is the medium to which we all, even the compulsive writers among us, naturally turn when an issue engages our emotions. Above all, I think, speech is the natural medium of decisionmaking. In my own family, when we are doing something hard like

picking a college for a high-schooler, we will write some transactional letters, and we will do some reading to gain information. But the essence of the decision process is oral. If this is true for someone like me—someone whose whole career has been involved with writing—how much truer must it be for those who do not see writing as a way of life. This latter group obviously includes most of humanity; it specifically includes nearly all of those for whom our written products are intended. I do not conclude from this that the directorate should get out of the writing business. I do not think we can do our work in any other medium. But if it can be shown that the cognitive processes involved in writing and speech differ in important respects, then I think it follows that the relationship between our writing-based work and our speech-oriented consumers deserves close attention.

Here are some of the ways in which writing and speech differ:

- With speech, much of the communication takes place in ways that do *not* involve words: in gesture, in tone of voice, in the larger context surrounding the exchange. Speech is a complex audio-visual event, and the implications we draw—the chunks we form—are derived from a whole network of signals. With writing there is nothing but the words on the paper. The result may be as rich as with speech—nobody would accuse a Shakespeare sonnet of lacking richness—but the resources used are far narrower.
- Writing calls for a sharper focus of attention on the part of both producer and receiver. When you and I are conversing, we both can attend to several other things—watching the passing crowd, worrying about some aspect of work, waving away a passing insect—and still keep the thread of our discourse. If I am writing or reading I must concentrate on the text; these other activities are likely to register as distractions.
- The pace and pattern of chunking is very different in the two modes. With speech, one word or phrase quickly supersedes the last, and the listener cannot stop to ponder any of them. What he ponders is the chunk he forms from his perception of everything the speaker is saying, and he is not likely to ponder even that very intensively. He does have the opportunity to ask the speaker about what he has heard (an opportunity almost never available to a reader), but he rarely does so; the spoken medium has enormous forward momentum. In compensation, speech uses a much narrower set of verbal formulae than writing. It relies heavily on extralinguistic cues, and by and large it is more closely tied to a larger context that helps keep the participants from straying too far from a common understanding. In the written medium, by contrast, the reader can chunk more or less at his own pace. He can always recheck his conclusion against the text, but he has little recourse beyond that. All the signals a writer can hope to send must be in the written words.
- A reader is dealing with a finished product: the production process has been essentially private. A listener is participating in a transaction that is still in progress, a transaction that is quintessentially social.
- Partly because of the factors listed so far, writing is capable of more breadth and more precision than speech. Neither complex ideas nor complex organizations would be possible without writing. My own impression is that even in this television-dominated era, people attach more solidity and permanence to something written than to something spoken. Perhaps we have an ingrained sense that the products of speech are more ephemeral than the products of writing. But to achieve this aura of permanence writing sacrifices a sense of immediacy. A writer tends to speak with the voice of an observer, not a participant.

Thus it appears that a person working with speech is doing markedly different things than a person working with writing. This is true whether the person is acting as a producer or a consumer. Further, because everyone works at times with speech—whereas not everyone works in any

comprehensive sense with writing—speech is a far more general medium of exchange. I suspect that most people tend to transfer what they have learned from writing into the spoken mode: usually that is where information gets used.

The transfer from writing to speech is made easier because the chunks I form after reading your text are most unlikely to take in all the detail you have laid out. The only way to work back toward your degree of detail is to reread your text, and I will not take in everything even if I do so (much of what you want to say is not even set down explicitly). If I do not go back I will operate with the chunks as they currently exist in my mind (and as they combine with other chunks, at least equally powerful, that my mind was working with before). This does not necessarily mean that my new set of chunks will be either less complex or less valid than yours; it does suggest that what I glean from your prose will be altered and simplified. Because of the simplification, what I have gleaned will be easily available for me to use in the speech mode. I will be all the more inclined to use it in this way if I am oriented toward decision and action, either generally or in this particular situation. If such is my inclination, I probably will do just what I did when the issue of sending children to college arose: view the text mainly in terms of what seems useful and ignore the rest.

The distinction I am making between writing-based and speech-based cognitive processes is illumined by Walker Percy's distinction between news and knowledge. To introduce his thesis Percy asks us to imagine an isolated island inhabited by people with a well-developed culture. On the shores of this island arrive thousands of sealed bottles, each bottle containing a single assertion like the following:

- $E=MC^2$ .
- A war party is approaching from Bora Bora.
- The dream symbol, house and balcony, usually represents a woman.
- Being comprises essence and existence.

- In 1943 the Russians murdered 10,000 Polish officers in the Katyn Forest.
- Tears, idle tears, I know not what they mean.
- There is fresh water in the next cove.
- Chicago is on the Hudson River or Chicago is not on the Hudson River.

An islander, Percy avers, might experiment with various ways of organizing these messages, but in the end would put each of them into one of two categories. The first category would include all the scientific and formal statements, all the generalizations, and also all the poetry and art. Producers of such statements are alike in their

*withdrawal from the ordinary affairs of life to university, laboratory, studio, mountain eyrie, where they write sentences to which other men assent (or refuse to assent), saying, "Yes, this is indeed how things are."*

The second category would include statements that are significant

*precisely insofar as the reader is caught up in the affairs and in the life of the island and insofar as he has not withdrawn into laboratory or seminar room.*

The statements about the Bora Bora war party and the water in the next cove would be obvious candidates for this category. The categories are neither hermetic nor unchanging. The statement about the Katyn massacre might be in either. The first category Percy calls knowledge; the second, news.<sup>a</sup>

Percy continues that not only are there two categories of information, they are read from two different postures, there are two different kinds of verifying procedures, and there are two different kinds of response:

- *Nature of the sentence.* Knowledge can in theory be arrived at "anywhere by anyone and

---

a. Percy, *The Message in the Bottle*, 119–39.

at any time”; news involves a nonrecurring event or state of affairs which bears on the life of the recipient.

- *Posture of the reader.* The reader of a piece of knowledge stands “outside and over against the world;” the reader of a piece of news is receiving information relevant to his specific situation.
- *Scale of Evaluation.* We judge knowledge according to the degree it achieves generality; we judge news according to its relevance to our own predicament.
- *Canons of acceptance.* We verify knowledge either experimentally or in light of past experience. News is “neither deducible, repeatable, nor otherwise confirmable *at the point of hearing.*” We react to it on the basis of its relevance to our predicament, the credentials of the newsbearer (according to Percy, “a piece of news requires that there be a newsbearer”), and its plausibility.
- *Response of the reader.* A person receiving a piece of knowledge will assent to it or reject it; a person receiving a piece of news will take action in line with his evaluation of the news. (And, I would add, the receiver of a piece of news is more immediately concerned than the reader of a piece of knowledge with the correctness of the information.)

Obviously, I am building toward an assertion that the DI tends to deal in knowledge and our customers are interested mainly in news; and furthermore that there are correlations between news and the cognitive processes involved in speech on the one hand, and between knowledge and the cognitive processes involved in writing on the other. Equally obviously, the reality is not that simple, and the correlations are not exact. Not only does it demean our consumers to imply that they have little interest in knowledge as Percy defines it, but the distinction between news and knowledge, never airtight, has become increasingly problematic over time. With the information explosion and the widespread acceptance of the notion that knowledge is power (or

to be more consistent with Percy’s terminology, the notion that control over news is a source of power), speech- and action-oriented people have become ever more eager to scan knowledge-purveying documents for their news value. Moreover, the scope and depth of what such individuals are expected to know has expanded greatly. In both senses the potential domain of news has grown and has come to include regions that might in the past have been left to knowledge.

The boundaries get fuzzed in other ways. If I have discovered a new way to look at a problem, the discovery is likely to have the feel of news for me. But I may decide I have to use a knowledge-based mode to explain it. Perhaps, if I come from a scholarly environment, I believe knowledge has higher status than news. Or perhaps I do not think I can do justice to my discovery in a news-based mode—and if news, like speech, has a limited capacity to cope with complexity and my discovery is complex, my belief may have some validity. The recipient probably will duly register the resulting product as knowledge and may miss the sense of discovery altogether.

Finally, the material the DI deals with usually can fit under either heading, as was the case with the sentence about the Katyn massacre. A statement about the boiling point of water clearly is knowledge, and a statement about a fire that has just broken out in my office clearly is news; a statement about a balance-of-payments problem or a festering insurgency or a restive legislature could be either and is probably both. None of the latter statements is likely to have a high level of generality, none is likely to be easily verifiable under the canons of acceptance applicable to knowledge, and all are likely to involve the news-type question, “What (if anything) should the United States do?”

I still would argue, however, that when all is said and done the analytic work of the directorate fits more naturally on the knowledge side of the line and is inseparable from writing; whereas our consumers tend to look for news and to be more

comfortable with speech. I think this is one way of stating a basic challenge facing the directorate.<sup>a</sup>

I find it impossible to avoid the conclusion that we will do our work better if we include elements of speech and news. I also think that if we produce something that sounds like speech, it will tend to sound like news as well. I am not saying we should make our writing read like spoken English; the canons of speech and writing differ too greatly for that. But I do think we should aim for prose that has a conversational ring. How one does this is largely a matter of individual style, but some suggestions may be possible. Speech specializes in short sentences, it shies away from complex constructions, it uses short words (especially Anglo-Saxon ones), and it is peppered with verbs. Modelling prose around these characteristics—not inflexibly, not burdened by the specious claim that the result will inevitably be Dick-and-Jane (or perhaps Hemingway) prose—will, I think, help writing sound a little more like speech.

But don't you thereby rob yourself of conceptual complexity? Is it possible to lay out difficult ideas in conversational prose? All I can say is, it's been done. Whatever one thinks of Robert Pirsig's *Zen and the Art of Motorcycle Maintenance* in terms of substance, it does present some tough philosophical concepts in a palatable enough way to make the best-seller lists. The same is true of Hofstadter's *Godel, Escher, Bach*. There is, in fact, a tradition of heavyweight conversational prose dating back as far as St. Paul and Plato. So at least in theory the idea is feasible. I would also urge all writers to ask themselves the following question: the idea may have been hard to clothe in words and get down on the page properly, but now that it is there, does

it seem quite so complicated? Perhaps so; on the other hand, perhaps not. If the latter, perhaps putting it into more conversational language would not be too difficult.

This approach would nonetheless pose several problems. What I have called heavyweight conversational prose differs markedly in style from that usually produced by the directorate, and I think the differences arise from factors other than the DI's stylistic conventions and variations in individual writing skill. The most notable difference is that all the writers I cited as examples make extensive use of the first and second persons. This is a natural tendency when one is being conversational; one almost inevitably tries to personify the newsbearer and the recipient. I found in writing this essay that giving in to the tendency has a liberating effect.

A writer can forgo the first person singular and still be conversational—fortunately for the directorate, since the message-bearer for the news we write is a collective labeled “CIA,” which can call itself “we” but never “I.” The tendency to personify would persist, however. Even without the first person singular, I suspect the directorate's prose would take on a more individualistic cast. Moreover, a person writing in this mode is likely to assume a more actively persuasive tone and to risk crossing the line between persuasion and argumentativeness. These considerations lead me to believe that a turn toward conversational prose would add to the stress of the review and (perhaps even more) the coordination processes. I also think it is worth asking whether the structure would really be willing to tolerate a higher individuality quotient in the papers it produces.

I do not believe any substantial move would be made toward a greater news/speech content without tradeoffs such as these. They probably could be managed, if only because such changes in institutional style would have to come gradually, with large amounts of satisficing at every stage. But I do think such issues would have to be anticipated.

---

a. I also believe it is one factor making for turbulence in the flow of the review process. A reviewer, alive to the desires of the consumer, is likely to be looking for something resembling news/speech, whereas the writer will worry that the nuances and precision that can be conveyed only through knowledge/writing will be lost. What is pedestrian or convoluted to one is detailed and circumstantial to the other.

## IV. Writing Schemes and Cognitive Overload

Only the original writer can impart a conversational tone to a draft, of course; it cannot be edited in later. Yet this essay has suggested that a writer is not likely to get his ideas articulated with the precision they require unless he works in a knowledge/writing mode, and that knowledge/writing is not inherently conversational. This implies that the writer would have to rework his prose even after the ideas had been articulated—a prospect few writers would relish. But putting the issue in these terms overstates the added burden and misstates the way it would be felt. The writing process is extraordinarily convoluted, and all writers constantly shift from one aspect of the problem to another. In this juggling act there is no question of deciding what you are going to say and then deciding how you are going to say it; writers—good writers in particular—work on all aspects of their problem virtually from beginning to end. If Flower and Hayes are correct, moreover, good writers represent this problem to themselves as a “complex speech act”; the conversational element is already present to some degree.

Looking at the writing process as a whole will illuminate these notions. A writer is trying to accomplish two quite different things: to define what his ideas are by clothing them in words, and to communicate those ideas to others. Such is the nature of writing that the two goals are inextricably intertwined. To repeat, the writer does not decide what his ideas are and then decide how to communicate them; internal and external communication are of a piece. The complexity of this operation is one source of “cognitive overload.” Another source is that as he writes, the writer is creating a datum—a malleable entity outside the mind that grows out of the mind’s internal workings. The datum acts as an extension of working memory; the writer now has available a dependable array of chunks that is not limited by what working memory can attend to. Yet working memory itself is as tiny as ever, and it can-

not keep track of all aspects of the datum unless it shifts constantly from one aspect to another.

Skilled writers use a variety of means to reduce the overload. For example, they satisfice. “Not the right word but what the hell,” said one subject of an oral-protocol experiment at Carnegie-Mellon. The context indicates he was implying that the word was good enough for now and could be polished when working memory had been cleared of other demands. Or they draw on patterns stored in long-term memory, just as the skilled chess player does. (The patterns must be appropriate, of course. A new arrival in the DI is likely to discover that many patterns learned earlier no longer work, and building a store of, say, 50,000 patterns that do work takes time.) According to Flower and Hayes, perhaps the most effective technique of all is what they call “planning”—but the term takes on a different meaning in this context. It does not simply equate with outlining. In fact, Flower and Hayes found in their analysis of oral protocols that few people made use of outlines or other structured techniques. “Planning” might better be thought of as developing heuristic strategies and monitoring those strategies.

Flower and Hayes suggest that there are three different sorts of strategies:

- At the highest level, a strategy *To Do*. This is the rhetorical problem the writer sets out for himself: “Write a current intelligence article shooting down this coup report”; “Turn out something that will get through the review process without too much hassle”; “Write a paper that shows how stupid the conventional wisdom is”; “Set down this new idea.”
- Next, a strategy *To Say*. This is essentially a content plan—the points to be gotten across to the reader. One can jot down informal notes or work up a detailed outline. “A plan *To Say* is

essentially a scale model of the final product. Perhaps that is why it has been so widely and vigorously taught, often to the exclusion of any other kind of planning.” The plan *To Say* is subordinate to the plan *To Do*.

- Finally, and coordinate with the strategy *To Say*, a strategy *To Compose*. This category includes the interaction between ideas and the developing text. It includes short-range mental notes like “I’ll write down a bunch of ideas and connect them later.” or “A point I will want to make someplace is that...”

Failure to go beyond a plan *To Say* is a good route to ineffective prose. Skilled writers are adept at using all three strategies in conjunction with each other, checking and rechecking one against another and monitoring how each is working. They spend less time considering plans *To Say* than do novices, and more time considering the overall assignment (plans *To Do*) and the rhetorical challenges involved (plans *To Compose*). Flower and Hayes add: “Moreover, as they write, they continue to develop their image of the reader, the situation, and their own goals with increasing detail and specificity.”

A striking aspect of the approach of skilled writers is the frequency with which they think about how they are affecting the reader. A comparison between two writers of different skill levels is shown in the table below. The expert spoke of the way he represented the audience and the assignment more than twice as often as the novice did, and he spoke of goals vis-a-vis the audience 11 times, whereas the novice did not consider this aspect of the problem at all. By contrast, the two were quite similar in the frequency with which they addressed questions of text and meaning. Hayes and Flower go on to assert that being alive to the audience and other aspects of the rhetorical context enriches

substantive content as well. With a good writer three fifths of the new ideas grew out of thoughts about the assignment, the audience, or the writer’s own goals; whereas with poor writers 70 percent of the new ideas flowed from the topic itself. “All of this suggests that setting up goals to affect a reader is not only a reasonable act, but a powerful strategy for generating new ideas and exploring...a topic.”

Here we have, I think, a way of closing the gap between news/speech and knowledge/writing. Perhaps, in setting up the dichotomy between news and knowledge, Percy underestimated the wiliness of the skilled writer. Perhaps the body of information that he calls knowledge is roughly equivalent to what a writer would work with under a plan *To Say*—the substantive points to be made. But Flower and Hayes say the mark of a good writer is the resources devoted to the aspects of the problem other than the substantive points, especially the strategy *To Do*. Attention to these areas—and in particular, treating the audience as a vivid entity—is what distinguishes those who can turn the process of creating text into a “complex speech act.” For such writers, conversational prose not only is possible, it is what they tend to produce.<sup>a</sup> Thus to move writing back toward speech and knowledge back toward news, a writer should be urged to treat the nonsubstantive aspects of the assignment with the importance they deserve. It usually will not be possible to retain all the density of argumentation that knowledge/writing can achieve, but a surprising percentage can be preserved, and what does get set out has a much better chance of actually being transferred.

a. This notion is buttressed by the examples of “good writing” included in my favorite how-to-write book, Jacques Barzun’s *Simple and Direct*. The writers whom Barzun singled out range from Dorothy Sayers to Eric Hoffer to a man writing about how to use a saw. Nearly all of them, and Barzun himself for that matter, produced what I have been calling conversational prose.

	Analysis of rhetorical situation: Audience and Assignment	Analysis of goals				
		Audience	Self	Text	Meaning	Total
Novice	7	0	0	3	7	17
Expert	18	11	1	3	9	42
Number of times writer explicitly represented each aspect of the rhetorical problem in first 60 lines of protocol						



## V. Creativity and the Conceptual Front End

One of the enduring concerns of DI analysts is the creativity issue. What role is there for creativity in a structured, basically hierarchical organization—an organization, moreover, that operates more or less according to craft-work standards? Doesn't such an organization run the risk of stunting creativity, making do with run-of-the-mill analysis, and perhaps missing important trends? On the other hand, some might argue that creativity is not even particularly relevant to what we do. The important goal is to meet the needs (stated and implicit) of our consumers in a timely and accurate way, and tons of creativity will not help if this goal is not met. An organization has both the ability and the obligation to impose norms, these people might argue; constraints of some sort are bound to turn up sooner or later, and they almost certainly will seem arbitrary when they do. So why not impose the constraints of the craft, which at least have the virtue of some consistency and continuity?

This argument is clouded by a definitional problem which needs to be disposed of before we can address the issue constructively. Creativity usually is treated as a rare phenomenon, and so it is under the usual definition. But the whole cognitive process actually has a creative component, and in this specific sense creativity is all around us. I believe many of the problems confronting an organization like the DI, whose *raison d'être* is the intellectual activity of the people who comprise it, flow from the conflict between these two sorts of creativity.

Herbert Simon provides a useful list of the criteria by which the creative aspect of a piece of work is judged:

- The work has novelty and value (either for the originator or for his culture).

- It is unconventional in the sense that it requires modification or rejection of previously accepted ideas.
- It requires high motivation and persistence, taking place either over a considerable span of time (continuously or intermittently) or at high intensity.
- The problem as initially posed was vague and ill-defined, so that part of the task was to formulate the problem itself.<sup>a</sup>

Simon notes that not all the criteria need be satisfied before a work is considered creative. I would go further and assert that for the individual producing the work, one criterion suffices by itself: that the work have “novelty and value...for the originator.” Furthermore, the nature of the

*What  
role is there for  
creativity in a structured,  
basically hierarchical  
organization  
?*

cognitive process is such that almost everything produced will meet this criterion. I noted early in this essay that a child's approach to the learning of a language has a creative as well as a satisficing element. A child discovers how language works by constructing original sentences, sentences that did not exist until he formed them. I suspect that pleasure at having produced something new is a powerful force for further linguistic exploration and that an overlearned (and thus unconscious) sense of satisfaction accounts for part of the conversational dynamic in adults.

The creative component in this narrow, individualistic sense is if anything even more apparent with written work. With speech, the words are ephemeral and part of a social process. The individual producing them cannot get too closely identified with them; they are tossed into the conversational stream and soon are lost to consciousness, leaving nothing behind but an assortment of conceptual chunks in the memories of those participating

---

a. Simon, *Models of Thought*, 139.

in the exchange. (People do, of course, identify themselves closely with the chunks.) By contrast, a person producing a written sentence is bringing something discrete and durable into existence, and I would suppose that the sense of having created something is all the stronger for that reason. The private nature of the writing process reinforces this sense. Moreover, the imbalance in emotional investment between producer and consumer—the one closely involved with the specific words, the other much less so—is likely to be more apparent with writing.

The writer must circulate his text to others; this is part of his unspoken contract with the rest of the world. The criteria at this point, however, are the other items on Simon's list, which are all essentially social rather than individual, and society discerns creativity less often than the individual does in himself. An organization, moreover, is necessarily concerned with standards of some sort, and it has a right to require a degree of conformity from the journeyman.

Acceptance of such norms does not always hobble creativity. Until the last couple of centuries much of the great art was produced by people whose outlook was that of craftsmen. Consider the following comment by Jerome Bruner:

*One cannot help but compare the autobiographic fragment left by Ghiberti, discussing the long period during which he worked on the famous doors of the Baptistery at Florence, with the personal writing, say, of a modern sculptor like Henry Moore. Ghiberti talks of the material that was "needed" to do the designs that were "required." It is as if it were all "out there." Moore is concerned with the creating of illusions and symbols, and self-awareness for him is as important as a stone chisel.<sup>a</sup>*

Yet it does seem to me that creativity is important for an organization like the DI, and that such an organization constantly runs the risk of inhibiting the creative component in the work of its people. Few, I suspect, would dispute the proposition that

satisfying consumer needs is aided by a certain degree of creativity, and to me it is equally obvious that some sort of creativity is essential if the directorate is to satisfy the predictive aspect of its mission as well as possible. I would also suppose that work deemed creative under Simon's broader, socially determined criteria must have creativity in the narrow sense as a precondition. It is here that the risks arise for an organization like the DI. The following examples cited by Morton Hunt support this argument.

*In one of Torrance's many studies of creativity in school children, children were given pencils, crayons, and simple collage materials and told to make a picture that no one else would be likely to come up with. A variety of kinds of comment and appraisal were given to different groups of the children, and some got no evaluation at all. The upshot: those who worked without evaluation turned out pictures that were more creative than those produced by children receiving the most constructive commentary. In a study by another researcher, a group of students were told they would earn a reward for thinking up the largest number of plot titles and stories; their output was less imaginative and original than that of another group who expected no reward. Merely, knowing that one's work will be critically appraised by experts after it is finished has a negative influence. Teresa Amabile, in her collage studies, told one group of college women that their efforts would be judged by artists, while another group was told nothing; the latter group did significantly more creative work.<sup>b</sup>*

One can easily criticize research such as this. For example, how is it possible to judge something as subjective as creativity? Might not the researchers have gone into their studies with biases that skewed the results? Perhaps a more telling point is that the subjects of the experiments seem not to have had much skill or experience in the activities they were engaged in, and thus may have been particularly sensitive to the notion of being judged. One might expect a more complex situation, for

---

a. Bruner, *On Knowing*, 54.

---

b. Morton Hunt, *The Universe Within*, 313.

instance, if an experienced artist was faced with the possibility of being judged by other artists. But even taking such objections into account, the notion of creativity being inhibited by the mere prospect of outside comment—even positive comment—is thought-provoking. Even a skilled artist might be inhibited if he was required to keep in the front of his mind the evaluation to which his work was to be submitted. Yet this is a requirement which a structured organization can hardly avoid imposing.

Or consider the comment by Jacob Bronowski in his Silliman lectures at Yale in 1967:

*The society of scientists, the community of scientists, has this advantage, that from the moment we enter it, we all know that fifty years from now, most of the things we learned here will turn out not to have been quite right. And yet that will have been achieved without enormous personal dramas. It will be achieved by giving due honor to the people who take the steps, the steps that turn out to be wrong as well as the steps that turn out to be right.<sup>a</sup>*

Again one can quibble: surely it is only the “steps that turn out to be right” which win Nobel prizes. But the essence of Bronowski’s argument is hard to dispute. Science does reserve a place of honor for those who explore blind alleys, not least because discovery is as likely to proceed from an earlier error or ambiguity as from an earlier truth. The place of error in a structured organization, on the other hand, is much less certain, particularly if accountability is one of the driving forces in that organization. Even with the best will in the world such an organization is likely to put a premium on being right. As a result tentative ideas may have a harder time surviving, and there may be a tendency to equate “being right” with “not being wrong.” If so, caution and the school solution will have an easier time of it than they might otherwise have.

And yet the Directorate of Intelligence cannot devote many of its resources to nurturing creativity, at least under any very expansive definition of the term. Like any craft organization, the DI has jobs to do, and worrying about creativity would often get in the way. Moreover, much of the activity we report and analyze is actually pretty ordinary or closely tied to a narrow set of events, and with work of such an ephemeral nature the scope for creativity is limited. Finally, there is no guarantee that you will get more creativity if you do nurture it. I know of no approach that has proven more capable than others of producing results which are simultaneously creative, valid, relevant, and efficient.

*Yet, the  
DI cannot devote many  
resources to nurturing creativity... it has  
jobs to do and worrying about creativity would  
often get in the way.*

The concepts discussed in this essay do, however, offer some hints about reducing the tension between structure and creativity—reducing rather than eliminating, since the sources of tension will not go away and we are certain to wind up satisficing. The key is the importance of the work we do right at the outset of the process. Simon notes at one point, “Much of the skilled processing in chess occurs at the perceptual front end.” This proposition holds for all cognitive activity, although in the DI’s case it may be more accurate to talk about a conceptual rather than a perceptual front end. Being clear what is going on at the beginning takes on overwhelming importance in light of cognitive scientists’ findings; moreover, there is a built-in potential for conflict between what happens at the conceptual front end and what happens during a necessarily serial review process.

The importance, as well as the difficulty, of defining what is happening at the start is put into clearer relief by an obvious point made by Flower and Hayes: “People only solve the problems they give themselves.” I cannot give you a problem to solve; I can try to interest you in a problem; I can talk it through it with you; if I am your boss, I can order you to take it on. But the problem you solve still will be the one you pose to yourself, not the one I have given you. The two will rarely be congruent, and if we are not clear with each other they may diverge

---

a. Jacob Bronowski, *The Origins of Knowledge and Imagination* (New Haven: Yale University Press, 1978), 133.

drastically. Each of us will have sharply narrowed our concept of the problem by the time we begin to talk about it, and neither my concept nor yours will be confined to the substance of the issue. We both will have a complex set of assumptions as well as an assortment of goals, usually unspoken, that we hope to achieve. If I am your boss one of my goals might be, “Try once more to get some worthwhile prose out of this analyst.” You, the subordinate, might have “Get this guy off my back” as one of yours. Both goals will affect the product—your goals more than mine, in fact, since as the one in charge of the keyboard you have the biggest role in deciding how the problem will actually be solved.<sup>a</sup>

The problem you give yourself to solve is roughly equivalent to your strategy *To Do*—the rhetorical challenge in all its complexity, as it is posed by the person meeting it. This is easily the most important part of the “conceptual front end” to articulate, particularly if the topic is difficult or controversial or if it is breaking new analytic ground. I suspect it is the part that is least well articulated for much of the DI’s work. Most of the time we focus on the areas of *To Say* and *To Compose*—the points we are going to make and the way we are going to express them. But problems in these areas are often easy to fix. It is the unspoken divergences at the level of plans *To Do*, I believe, that cause the real headaches. This is particularly true if such divergences do not come to light until a paper makes its way up through the hierarchy.

It is terribly hard to articulate one’s strategies at this level. In the ivory-tower setting of this paper I can make the obvious point that everyone benefits when a goal like “Show those turkeys downtown how stupid they are” is brought into the open, or when a reviewer-to-be makes explicit a line of attack that looks promising. In a real-life situation, however, it is in the nature of assumptions that they do not even get noticed, much less articulated. And articulated in the context of the specific situ-

ation: if an assumption is not so articulated, you can never be sure it is common property, given the idiosyncratic way each of us chunks the information at our disposal. The post-facto “any-fool-should-have-known” argument is always an unsatisfactory substitute.

A sure way not to illumine a strategy *To Do* is to develop nothing but an outline, since outlines are necessarily concerned mainly with plans *To Say*. Concept papers will shed more light, but in my experience they rarely are fully satisfactory. This is at least partly because strategies *To Do* typically get articulated bit by bit (once again we meet the proposition that learning and insight are achieved heuristically and incrementally). Thus the actual process of finding out what has happened at the beginning of the conceptual process takes time. It involves a series of exchanges, not just at the formal inception of a project but even before it takes shape and also as it progresses.

Not all analysts find such exchanges congenial, managers at all levels begrudge the time, and divergences among various strategies *To Do* are not always resolved amicably. And indeed not all projects are worth spending too much time on. But to the degree that all concerned—higher level reviewers as well as those actually involved in the writing—can work early and often at narrowing the gaps among various conceptions of the rhetorical problem, the chances of a bumpy review process will diminish. Perhaps more important, the chances of preserving whatever creativity there is will be enhanced. And finally, hammering out differences at the level of strategies *To Do* might make a virtue of necessity: it might give a more collegial cast to a process whose hierarchical aspects are built-in and inescapable.<sup>b</sup>

---

b. Looking at the problem in this way exposes an anomaly in a related area. As we have seen, the value of experience—experience extending over a decade or more—is well documented in the cognitive-science literature. Such experience generally makes itself felt at higher levels of problem solving; a first-class composer or a chess grandmaster stands out because of his skill at a level analogous to Flower and Hayes’s strategy *To Do*. In the DI the people in upper levels of management have a high concentration of experience of this sort, yet the serial nature of the review process and the other demands of their job make it hard for them to bring it to bear at the point where

---

a. Many analysts may cavil at the last statement, arguing that the review process has left them no longer in charge of their own prose. Perhaps analysts do have less control than in earlier eras, but I would argue that because they are the ones constructing the prose, for better or for worse they remain by far the most influential factors in what gets produced.

*(Editor's note: At the time of this monograph's original appearance there were developments on the horizon that promised profound effects on the way DI analysts did business in just these areas. Project SAFE and the concurrent development of AIM appeared with recurrent troubles that tended to buttress the negative mindsets of many in the directorate at the time. Their potential attractiveness became apparent before long. The attractiveness in this context arose not from the information-retrieval aspect of SAFE, but from the interactive capabilities of the partner development, AIM, which was designed to permit frequent written exchanges between analysts—December 2009.)*

If the history of such systems is any guide, DI analysts are likely to find interaction of this sort a congenial way to work. If they perceive the initial capability to be inferior to what is available elsewhere, they are likely to press for greater speed, comprehensiveness, and flexibility.

AIM's interactive capacity will be a boon in many ways, but it will also pose new managerial challenges:

- It will encourage collegiality among analysts concerned with various aspects of the same problem—an important advantage in the expanded directorate. It will not magically eliminate parochialism and dog-in-the-manger attitudes, but to the extent it is used, it will help analysts establish communication and get around competitive barriers.
- Exchanges via AIM will be informal and basically conversational, and there is every reason to expect some of the conversational flavor

to carry over into more formal products. Thus AIM seems likely to foster the sort of prose called for earlier in this essay. If it does, of course, it will also raise the sort of managerial questions discussed there.

- AIM seems well suited to the incremental articulation of strategies *To Do*. Those involved in the exchanges are likely to develop clearer notions—and a higher degree of consensus—about what their purposes are.

*There  
never will be  
dependable correlation  
between the intensity of the interaction and  
the quality of the product.*

- But those who do not participate will have a harder time catching up. This means, among other things, that potential reviewers may find it even more difficult to make contact with the problem-defining process unless they keep current as it progresses—unless they participate in the written as well as the spoken interplay. The potential for heavier demands on managers' time is obvious.

- At the same time, collegiality could work at cross purposes with hierarchy. What would it do for a division or branch chief's authority, for example, if the office director used AIM to get deeply involved in helping an analyst define a rhetorical problem?
- Nor will the system be uniformly beneficial on the substantive side. Writing still will be essentially a private process, and some analysts will not function at their best if they have to work in a sort of electronic marketplace of ideas. Others may find the exchanges so much fun that they forgo analytic digging. There never will be dependable correlation between the intensity of the interaction and the quality of the product.

---

strategies *To Do* are evolving.

---

## VI. What Next?

Out of this discussion grow numerous questions about the way the directorate goes about its business. They include the following subject areas.

### 1. Recruitment.

Is there any way in pre-employment interviews to focus more sharply on the way a prospective employee deals with strategies *To Do* (and perhaps *To Compose*)? More pointedly, in hiring decisions should more weight be given to these attributes and less to skills of the *To-Say* variety? Should we look harder for writers and reporters and less hard for scholars? Should we count more on internal training and put less reliance on what an individual has learned before coming aboard? Should there be sharper distinctions in approach between disciplines—would it be reasonable, just for example, to require advanced degrees for economists but to emphasize BAs among political analysts? There never will be absolute answers to questions like these, but the ideas laid out in this essay make affirmative answers plausible in many cases. It also would be worthwhile seeing if there are any tests that get at skills of the *To-Do* variety.

### 2. Training.

If it is correct that around 10 years are needed to acquire the conceptual network necessary for first-class work, what are the implications for the way the directorate goes about its training effort? Presumably an analyst is well into the notional 10 years on his or her arrival, but it still might be prudent at that point to anticipate another 5 years of development. If so, it seems to me, more careful attention needs to be given to how the analyst learns the trade, and training should figure far more heavily in the DI's thinking than it does now. I would suggest that if what is involved is, in fact, craft-work, nearly all the training should be on the job. It even might make sense to make training one of the specific functions discussed on a branch chief's performance appraisal report (PAR). The question would then shift to how branch chiefs should be

trained. (An exercise of this sort might also produce a broader examination of the branch chief's functions—an examination I think the directorate would find illuminating.)

### 3. Analytic Writing.

The suggestions in this essay about conversational prose mesh with some of the trends already underway in the DI. If the best route to that sort of writing is to build a vivid image of one's audience, then frequent contact with the audience is likely to prove a useful tool for sharpening the writer's effectiveness. Moreover, the notion of bringing the audience into focus should help writer and reviewers establish common ground to work out the nonsubstantive aspects of a paper.<sup>a</sup> Beyond that, I would suggest that the vocabulary of cognitive science (and the concepts behind it) will be useful to those involved in teaching analytic writing, both on the job and elsewhere.

### 4. Organization.

I have suggested that there is an unavoidable conflict between the way an organization operates and the way the individuals in the organization do their cognitive work. Estimating the costs imposed on the organization by this conflict is impossible, since to do so one would have to guess what might have been produced but wasn't. It seems safe to say, however, that although the costs probably are lower than analysts believe, they probably are higher than many managers would admit. If so, some effort to ameliorate the conflict might prove worthwhile. Four lines of approach come to mind.

---

a. To an analyst, of course, the audience includes the chain of review; this may actually be the most important audience of all from his or her perspective. In addition, it should be obvious that an analyst will get a different sort of audience-image from working-level colleagues than from exposure to their bosses, and that secondhand information about the interests of the latter will have a tough time competing with the first-person experience. This is one of many conceptual gaps it will take repeated effort to bridge.

- *First, divorcing substance from hierarchy as often as possible.* This means exchanges between managers and analysts at times when a paper is neither under review nor in prospect—at conferences, for example, or simply as part of everyday chitchat. The greater the hierarchical gap, the harder such exchanges are to arrange, but the more valuable they might prove to be. To the extent that nonhierarchical channels for feedback can be developed, I believe the chances of operating at cross purposes when something is on the line will be reduced.
- *Second, hierarchy-jumping in contacts with consumers.* The desk analyst, in my opinion, would benefit enormously from learning first-hand what is on the minds of high-ranking officials.
- *Third, enfolding AIM's interactive function into the culture of the directorate.* Interactive systems are basically democratic, and I believe that properly used, they can reduce the height of organizational barriers.
- *Fourth, giving more articulation to assumptions at the level of strategies To Do.* One way this might be done is to modify the format of concept papers. Putting such papers in the first person, and replacing impersonal constructions ("This paper will examine...") with straightforward questions ("What will happen if...?") would help to make the rhetorical scheme more visible. AIM and extrahierarchical exchanges will have a similar effect.
- If a videotape were to be produced in which an analyst discussed the findings of a freshly produced paper, those findings would almost certainly make it across the gap between knowledge and news. Many in the DI's potential audience would probably find an extemporaneous discussion much more vivid than words on paper. The list of objections to the idea, of course, is formidable—the loss of precision as the analyst speaks extemporaneously, the time involved in producing even an informal discussion, the problem of defining the audience, the dilution of the sense that the paper is a product of the directorate as an institutional entity. But the benefits, if the tape was done right, might be equally impressive.
- Television might offer a better way of laying out alternative hypotheses. I have trouble believing that the present method, in which heated disagreements are set forth in determinedly flat prose, captures enough of the flavor to be much help. Taping a live discussion among experts might do a better job.
- In fast-breaking situations of prime importance, it might be useful to supplement other forms of current intelligence with televised analytic commentary. These days, current intelligence at its best almost exactly fulfills Walker Percy's criteria for news: it produces information bearing on the reader's predicament and it is carried by a newsbearer. Videotape would let a knowledgeable analyst be called into service as a specialized newsbearer on matters of great moment.

### 5. Presentational Methods.

Conversational prose is one way to bridge the gap between speech and writing, but are there others? Can more extensive use be made, for example, of television, whose effectiveness as an intelligence medium has already been proven in a limited way? I am thinking of something less formal, more ephemeral, and more personalized than what has been produced so far. Three possibilities spring to mind.

### 6. Further Exploration.

The DI might benefit from using cognitive-science techniques to analyze the production process. After reading an earlier draft of this essay, for example, Professor J.R. Hayes of Carnegie-Mellon University suggested a "cognitive task analysis of what analysts do. This information could aid the design of systems such as SAFE [and AIM]." Hayes has also suggested that we could do our own oral protocols of an analyst—or a reviewer—at work; he adds that

observing the reviewer might give particularly illuminating insights into the process.

Moreover, the notions in this essay exploit only a part of the expertise in cognitive science. For example, there is considerable disagreement in the field over the role of language in the cognitive process. Many would argue against the emphasis I have placed on the written word. They would say

that such concentration significantly limits the directorate's analytic flexibility. Others might assert that I have understated the degree to which computers can supplement our mental processes. Exploration of either field (they are not mutually exclusive) might bring further benefits to the DI's analytic effort. At this point, however, we leave the area in which I can even pretend to competence.

### Afterword

This list of suggestions points up once again the difficulty of the cognitive activities in which the DI is engaged. The difficulties arise because, as this essay has tried to show, the directorate must mediate a series of irreconcilable demands:

- Between the individual, private nature of the analytic effort and the social and hierarchical constraints imposed by bureaucratic imperatives.
- Between the cognitive patterns that are necessary to develop complex ideas (patterns that can only operate through writing) and the patterns our speech- and news-oriented customers are comfortable with.
- Between the early point at which critical analytic decisions are made and the serial nature of the review process.
- Between the need to search for new insights—to explore the murkier reaches of a problem space—and the need to avoid mistakes.

These conflicts have always been present, and we probably have always been aware of them, more or less. But as the DI grows and works out the implications of the trend toward centralization, as it spreads its analytic net more broadly and uses a finer mesh, they are likely to become more acute.

It would be easy to find a catalogue of conflicts like this depressing—to feel a bit like the centipede after its mind was set churning by the frog. But un-

like centipedes (even sentient ones), humans can also treat such conflicts as a challenge. Indeed it seems to be part of our nature to see a situation as a problem to be solved. Morton Hunt notes:

*We perceive situations as problems, and therefore undertake to solve them. An ape, coming to a broad river, would see it merely as an end to further travel in that direction; a human being might see it as a body of water to be crossed, and thereupon invent a raft. Innate neural impulses and early learning provide each species of animal with the specific procedures it needs to obtain food, avoid enemies, mate, care for its young, and so on; human beings, too, acquire procedures for dealing with these basic problems, but they also solve countless others that did not exist until their own minds saw them as problems. Art and arithmetic, music and money, detergents and democracy are all solutions to problems created not by Nature but by the human mind.*

*Problem solving is virtually species-specific, but what is absolutely species-specific and ultimately human is problem generating. The problem is in the eye of the beholder, and we are beholders.<sup>a</sup>*

We are, in fact, better at generating problems than at solving them. Rarely can we say unequivocally that we have solved a problem, particularly if it is a hard one. Rather we keep plugging away

---

a. Hunt, *The Universe Within*, 240.



at it heuristically. Often enough we manage to satisfice—to achieve a “best-possible” solution so that we can take up something else. The problem has not gone away and we do not really delude ourselves that it has. But being realists (i.e., heuristic individuals to the core), we accept the results and live with them—recognizing that sometime we may find ourselves returning to the problem once more. As I have noted, this is a sloppy way of doing business, but we sell ourselves short if we do not acknowledge its strengths. It deals as well as anything yet developed with the fact, to use Jerome Bruner’s words, that our “ability to deal with knowledge is hugely exceeded by the potential knowledge contained in man’s environment.”

The gross imbalance between what humans can know and what is available to be known is nowhere more apparent than in the Directorate of

Intelligence. Indeed this is the central dilemma the DI faces. We have always coped with the dilemma in a typically heuristic, satisficing way, and there is no reason to suppose this approach will change. My own impression, however, is that the directorate has never articulated the nature of the dilemma very clearly—a fine example of how hard it is to talk about the “conceptual front end.” A clearer notion of what the pressures are, and of what there is in our own makeup that makes them so powerful, should enable us to make better use of the power inherent in the heuristic approach. We never will get away from the need to satisfice—in other words, we always will be able to look at a given solution to a problem and see ways to make it better—but with greater understanding of what we are about, we may be able to attain a level of satisficing that we are more content with.